The economics of constant function market makers

Michele Fabi, Julien Prat

The Economics of Constant Function Market Makers*

Michele Fabi† Julien Prat‡

September 2024

Abstract

We use microeconomic theory to describe the inner workings of Constant Function Market Makers (CFMMs). We show that standard results from consumer theory apply in this new context, endowing us with powerful tools to characterize the optimal design of CFMMs. We employ them to analyze the externalities that traders and liquidity providers exert on each other when interacting through a CFMM. Liquidity providers reduce the execution costs by flattening the bonding curve on which trades are executed. Arbitrageurs impose an adverse selection cost on liquidity providers by unfavorably rebalancing their portfolio. We show that the strengths of these two externalities are pinned down by the curvature of the bonding curve and are inversely related to each other, thereby identifying the fundamental economic tradeoff that market designers have to address.

Keywords: Automated Market Makers, Decentralized Finance, Blockchain, Market Design.
JEL Code: D47, D53.

1 Introduction

Decentralized exchanges enable traders to transact without relinquishing the control of their assets to a third party. Instead of delegating the execution of their orders to a financial institution, traders can now rely on smart-contracts. Disintermediation removes the need for trust and thus holds the promise of lower transaction costs. But this benefit is counterbalanced by the operational costs of processing transactions through smart-contracts. The

*We are thankful to Fayçal Drissi, Weijia Wang, and Jan Christoph Schlegel for valuable discussions.
†Telecom Paris, CREST, IP Paris (michele.fabi@ensae.fr)
‡CNRS, Ecole Polytechnique, CREST, IP Paris (julien.prat@ensae.fr)

dominant design of centralized markets, namely limit order books, is not yet a viable option because it is too cumbersome to be efficiently handled in a decentralized environment. This is why leaner protocols called Constant Function Market Makers (CFMMs hereafter) have risen to prominence, establishing themselves as the main paradigm for decentralized exchanges.

The traditional approach to foster liquidity relies on professional market makers that stand ready to process incoming market orders. Automated Market Makers (AMMs hereafter) replace human market makers with algorithms. CFMMs are a subclass of AMMs whose minimalist design aims at achieving computational efficiency. They are built around liquidity pools. The level of inventory or reserves are determined by liquidity providers who can fill or deplete the pool. New liquidity provisions increase the size of the pool but leave constant the share of each asset in reserve. Traders, on the other hand, exchange one asset against another and thus alter the composition of the pool. The rate at which the exchange is processed corresponds to the market price quoted by the AMM. The main design choice therefore consists in determining the relative price of each asset as a function of the composition of the liquidity pool. A practical solution consists in defining an arbitrary *trading function* which depends solely on the reserves, and to ensure that the AMM only accepts trades that leave the value of the trading function unchanged, hence the name Constant Function Market Maker.

This simple design meets the two main requirements that a decentralized AMM should fulfill. First, it fosters market liquidity by guaranteeing that the AMM always quotes a price, and thus stands ready to process all incoming buy and sell orders. Second, the computational costs are low because the price only depends on an internal state of the smart-contract, namely the composition of its reserves. However, besides their practical convenience, little is known about the properties of CFMMs. What types of market structure and price discovery do CFMMs generate? More fundamentally, how should CFMMs be fine tuned so as to maximize the welfare of liquidity providers and traders?

One cannot address these questions without first identifying the tradeoffs involved in the choice of the trading function. We show in this paper that standard microeconomic theory sheds a surprisingly powerful light on this issue and, more generally, on the overall design space of CFMMs.

Considering first the problem of liquidity traders, we establish that it is isomorphic to the derivation of Hicksian demand. This fundamental insight enables us to apply the whole apparatus of consumer theory to the study of CFMMs. For instance, we show that the classical result according to which expansion paths are linear in wealth solely when the utility function is homothetic implies that the prices quoted by a CFMM are independent of

its pool size solely when its trading function is homothetic. This theorem drastically reduces the size of the design space since it suggests that, for most practical applications, designers can focus on trading functions that are homogenous of degree one.

After having characterized the optimization problem of traders, we turn our attention to its dual formulation. Again, we leverage similarities with consumer theory. We demonstrate that the minimal value of the portfolio held by liquidity providers can be derived following the same steps as the ones involved in the derivation of the expenditure function of consumers. This finding indicates that the rewards of liquidity providers are given by the solution to the dual problem. In other words, the two-sided nature of CFMMs is reflected in the structure of their optimization problems, with the primal capturing the perspective of traders and the dual that of liquidity providers.

We illustrate this general insight through two results of practical relevance. First, we show that *impermanent losses*, a concept widely used by liquidity providers to measure their exposure to adverse selection, are naturally expressed in the dual space because they are encapsulated in the expenditure function. Second, we prove that the trading set is fully described by the conjugate of the expenditure function. This theorem provides an intuitive interpretation for the otherwise opaque trading functions of CFMMs. They can now be directly constructed from the portfolio value of liquidity providers, making it possible to build CFMMs for sophisticated financial products.

We use duality theory to quantify and tie together the externalities that traders and liquidity providers exert on each other. We find that both externalities are a function of the trading function's curvature. A steeper trading function increases the marginal cost of trading but reduces the impermanent losses of liquidity providers. Moreover, the externalities exerted by traders and by liquidity providers are inversely proportional to each other.

To summarize, this paper demonstrates that the economics of CFMMs becomes apparent when examined through the lens of consumer theory. The similarities are striking as the unfolding of propositions closely follows that of microeconomic textbooks. Besides them, our research was also inspired by the seminal work of Angeris and Chitra (2020), Angeris et al. (2020) and Angeris et al. (2021a). They established some but not all the results presented in this paper. An essential difference is that they did so using a methodology rooted in convex analysis. Hence, our paper can partly be read as a translation of their research program into a language that is more accessible to economists. However, we believe that our contribution goes beyond pedagogical benefits since it unveils new economic intuitions and provides a unified framework for the analysis of CFMMs.

Related literature

We describe CFMMs as two-sided markets and provide a detailed analysis of their connection with microeconomic theory. Several recent contributions also analyze the economics of CFMMs but they do so from a different angle. Schlegel et al. (2022); Bichuch and Feinstein (2022) provide an axiomatic characterization of CFMMs. Schlegel et al. (2022) find that trading functions with constant elasticity of substitution fully identify the class of CFMMs that satisfy the axioms of independence and scale-invariance. Bichuch and Feinstein (2022) propose axioms that are satisfied by the vast majority of CFMMs. Bartoletti et al. (2021) provides another axiomatic characterization but from a computer science perspective. Jensen et al. (2021) analyze constant-product market makers, giving the intuition for some of the results that this paper proves in general. Park (2021) compares constant-product CFMMs with traditional market makers. He shows that constant-product CFMM are vulnerable to sandwich attacks, a form of miner extractable value. Jensen et al. (2021) also analyze constant-product market makers, giving the intuition for some of the results that this paper proves in general.

A growing stream of research studies the payoff profile of liquidity providers. Milionis et al. (2022) and Cartea et al. (2022) quantify the adverse selection cost of liquidity provision. They quantify impermanent losses resulting from a passive liquidity position and from a continuously-hedged liquidity position. Milionis et al. (2022) terms the latter 'loss-versus-rebalancing' (LvR). Cohen et al. (2023) determine the level of fees needed for liquidity providers to cover their losses and break even.

Cartea et al. (2022) also quantifies impermanent losses for concentrated liquidity CFMMs, i.e. with liquidity provision restricted to a specified price range. Bergault et al. (2022) outline a mean-variance analysis of the profitability of liquidity provision for a CFMM that incorporates external data through a price oracle. To the best of our knowledge, only Goyal et al. (2022) provide a characterization of optimal trading functions. They propose a convex optimization framework to design CFMMs that are optimal for a given specification of the price process of the traded assets.

Another relevant branch of literature studies the platform economy of decentralized exchanges. These papers propose variations of a benchmark static (or two-period) model to study the rents of traders and liquidity providers. Aoyagi (2020) determines the level of liquidity provision in competing CFMMs assuming both atomistic and strategic liquidity providers. Aoyagi and Ito (2021) cover instead the competition between a CFMM and a centralized exchange. Lehar and Parlour (2021) together with Capponi and Jia (2021) determine the impact of market fundamentals, such as noise versus informed trading, on the rents of liquidity providers. Using Uniswap and Sushiswap data, they both provide evidence

supporting their models.

Structure of the paper

Section 2 describes the functioning of CFMMs and explains why they create two-sided markets. Section 3 applies microeconomic theory to analyze the problem of arbitrageurs and its implications for the portfolio of liquidity providers. We also provide in this section the general conditions under which liquidity provision does not affect the prices quoted by a CFMM. Section 4 focuses on the welfare of traders and the externalities exerted by liquidity providers. Section 5 reverses the perspective of the previous section by focusing on the welfare of liquidity providers and the externalities exerted by traders. Section 6 uses duality to tie together both sides of the market. Section 7 introduces transactions fees while Section 8 discusses the objective of liquidity providers. Section 9 concludes while the proofs of the main results are relegated to the Appendices.

2 CFMMs as Two-Sided Markets

This section outlines the functioning of CFMMs and describes the market participants with whom they interact. A CFMM is a smart contract running on a blockchain. Each CFMM stores virtual assets' reserves in a *liquidity pool*. Based on these reserves, the CFMM performs two elementary operations: liquidity provision and asset swap.

CFMMs are platforms whose purpose is to connect *liquidity providers* with *liquidity consumers*. The two sides on the market exert cross-side network externalities on each other.[1] Liquidity providers (LPs) are the owner of the liquidity pool and constitute the supply side of the decentralized exchange. They provide the CFMM with its initial reserves and receive in return LP tokens which represent shares of the liquidity pool. Liquidity providers can at any time redeem their shares by burning (i.e. destroying) their LP tokens and receiving a corresponding fraction of the pool's current reserves. The demand side of the decentralized exchange is constituted of liquidity consumers. They interact with the CFMM by supplying some assets and withdrawing others from the pool's reserves. Thus we will simply refer to liquidity consumers as *traders*.

The behavior of each set of participants has a direct impact on the utility of the other set of participants. Liquidity providers benefit from the participation of traders because they collect fees that are proportional to the trading volume. Traders benefit from the

[1] For a definition of cross-side network externalities, see for example Tirole and Rochet (2003) or Armstrong and Wright (2007).

participation of liquidity providers as larger reserves translate into lower price impact. These network effects imply that the decentralized exchanges generated by CFMMs are *two-sided markets*.

The interactions between suppliers and consumers is mediated by a *trading function*, $U : \mathbb{R}_+^N \to \mathbb{R}$, that maps the reserve vector $R \in \mathbb{R}_+^N$ of the N assets held in the liquidity pool, with components R_i for $i \in \{1, \ldots, N\}$, to a real number $K \in \mathbb{R}$. The trading function encodes the set of trades that the CFMM will accept given its current reserves. The trading function represents the *trading set* in the same way that utility functions represent preferences in consumer theory. To underline this analogy, we use U to denote the trading function and refer to its value as the utility of the CFMM.[2]

2.1 Asset swap

Given reserves R, the CFMM accepts an asset swap that shifts reserves to R' if and only if $U(R') \geq U(R)$. The trading set is thus the trading function's upper-contour set

$$\mathcal{S}(K) \equiv \left\{ R' \in \mathbb{R}^N : U(R') \geq K \right\} \tag{1}$$

evaluated at $K = U(R)$.

It is convenient to decompose trades into input-output vectors $(I, O) \in \mathbb{R}_+^N \times \mathbb{R}_+^N$. $I \in \mathbb{R}_+^N$ is the vector of reserves that a trader inputs into the liquidity pool, whereas $O \in \mathbb{R}_+^N$ are the reserves that the CFMM outputs to the trader. The criterion for trade admissibility in Eq. (1) can then be stated in terms of netput (net output) vectors $\boldsymbol{\Delta} \equiv O - I \in \mathbb{R}^N$ as

$$U(R - \boldsymbol{\Delta}) \geq U(R). \tag{2}$$

Unless otherwise stated, we maintain a set of assumptions that guarantee that the trading function is well-behaved:

Assumption 1. *The trading function U is*

i Twice-continuously differentiable: $U \in \mathcal{C}^2$.

ii Strictly increasing: $\nabla U = (U_{R_1}, U_{R_2}, \ldots, U_{R_N}) >> 0$.[3]

[2] An alternative interpretation is to view CFMMs as firms. According to this analogy, trading functions are transformation functions from production theory. We favor the consumer analogy because it makes the analysis more transparent.

[3] Each gradient component is $\nabla U_i = U_{R_i} \equiv \partial U / \partial R_i > 0$. For $x, y \in \mathbb{R}^N$, the vector inequality $x >> y$ means $x_i > y_i$ for all i. Conversely, $x > y$ indicates that $x_i \geq y_i$ for all i and $x_i > y_i$ for at least one i. We assume that vectors are column vectors unless otherwise stated.

iii Strictly quasi-concave: If $U(R) \geq K$ *and* $U(R') \geq K$,

$$U\left(\alpha R + (1-\alpha)R'\right) > \min\left\{U(R), U(R')\right\} \quad \textit{for all } R \neq R',\ \alpha \in (0,1).$$

Assumption 1 summarizes standard regularity assumptions for utility functions in consumer theory. Differentiability (i) makes the trading function smooth, allowing for differential analysis. Monotonicity (ii) aligns the CFMM's utility with its reserves, so that the CFMMs accepts trades that either deepen its reserves or leave them invariant.[4] Quasi-concavity (iii) guarantees that $\mathcal{S}(K)$ is a strictly convex set. This implies that the CFMM prefers balanced reserves over extreme ones: If the CFMM accepts trades that leave reserves R and R', then it accepts a trade that leaves reserves at the average of those two. As we will see below, each of the three assumptions plays an important role in ensuring that a CFMM is well behaved.

Although the trading set includes all the trades which increase the utility of the CFMM, no rational trader would input more than required. This is why decentralized applications are configured so that traders leave the CFMM exactly indifferent between its pre- and post-trade reserves. We will therefore restrict our attention to trades occurring on the indifference curve reached by the initial reserves, i.e. on the *bonding curve* defined as

$$\mathcal{S}^b\left(U(R)\right) \equiv \left\{R' \in \mathbb{R}^N : U(R') = U(R)\right\}. \tag{3}$$

Quantity function and spot prices

We can use Eq. (3) to define the terms of an asset swap. Consider a trader submitting to a CFMM with reserves R a buy request of $\Delta \equiv \boldsymbol{\Delta}_i > 0$ units of asset i in exchange for asset j. The amount of asset-j input the trader has to pay is determined by the *quantity function*, $q_{ij}(\Delta, R) : \mathbb{R} \times \mathbb{R}^N_+ \to \mathbb{R}$, implicitly defined by:[5]

$$U\left(R^\Delta\right) - U(R) = 0, \tag{4}$$

where R^Δ are the post-trade reserves

$$R^\Delta \equiv R - \Delta\mathbf{i} + q_{ij}(\Delta, R)\mathbf{j} \tag{5}$$

with $\mathbf{i}, \mathbf{j}$ denoting basis vectors.[6] The quantity function can be differentiated to give the marginal price charged by the CFMM. Formally, we define the *price function* $p_{ij}(\Delta, R) \equiv$

[4] Monotonicity (ii) is equivalent to the path-deficiency property in Angeris and Chitra (2020).

[5] Notice that $q_{ij}(\Delta, R)$ is also defined for $\Delta < 0$, which indicates a request to *sell* asset i. In this case, $\left|q_{ij}(\Delta, R)\right|$ gives the amount of asset-j output that the trader will receive from the CFMM.

[6] $\mathbf{i} \in \mathbb{R}^N$ is such that $\mathbf{i}_i = 1$, $\mathbf{i}_{-i} = 0$. $\mathbf{j}$ is analogous. We will drop indexes ij when clear from the context.

$D_\Delta q_{ij}(\Delta, R)$, so that

$$q_{ij}(\Delta, R) = \int_0^\Delta p_{ij}(x, R)\,\mathrm{d}x. \tag{6}$$

To further connect with economic theory, a straightforward implicit differentiation of Eq. (4) shows that $p_{ij}(\Delta, R)$ is a marginal rate of substitution (MRS):[7]

$$p_{ij}(\Delta, R) = \left|\mathrm{MRS}_{ij}(R^\Delta)\right|, \tag{7}$$

where $\mathrm{MRS}_{ij}(R^\Delta) \equiv -U_{R_i}(R^\Delta)/U_{R_j}(R^\Delta)$ is the slope of the bonding curve measuring the amount of asset j the CFMM has to receive to compensate a marginal reduction in the reserves of asset i while maintaining its utility constant.

Multi-asset trade

A CFMM can also allow multi-asset trades. The most general way to formalize them with Eq. (4) is by replacing $\Delta\mathbf{i}$ with a generic output vector $O \in \mathbb{R}_+^N$, and $q_{ij}(\Delta, R)\mathbf{j}$ with a generic input vector $I \in \mathbb{R}_+^N$. However, doing so generally leads to a multiplicity of solutions for I given O. Thus multi-asset trade requires imposing further restrictions on the contract space in order to be well defined (Angeris et al., 2022). We handle multi-asset trade by considering *composite assets* defined by generic basis vectors $\mathbf{b}, \mathbf{b'}$ that replace $\mathbf{i}, \mathbf{j}$ in Eq. (5). With this interpretation, the analysis of two-assets trades naturally extends to higher dimensions, as we explain in Section 4.2.

2.2 Liquidity provision

The operations performed by liquidity providers can also be represented as input-output vectors. Liquidity provision is an operation with $I > \mathbf{0}$ and $O = \mathbf{0}$; while liquidity withdrawal has $I = \mathbf{0}$ and $O > \mathbf{0}$.

Most CFMMs advise or require liquidity providers to supply all the N assets simultaneously so as to maintain their shares in the pool. As we will see in Section 3, under common assumptions on the trading function this type of liquidity provision does not alter CFMM prices and therefore does not generate an arbitrage opportunity. More formally, the liquidity input has to move reserves from R to $R' = \alpha R \geq R$, where the scalar $\alpha \geq 1$ determines the CFMM utility at the new reserves. Thus the input vector has to be formatted as $I = R(\alpha - 1)$. In reward for supplying liquidity, the provider receives an amount of newly minted liquidity

[7] The identity is proven by totally differentiating U along the bonding curve to obtain $U_{R_i}(R^\Delta)\,\mathrm{d}\boldsymbol{\Delta}_i + U_{R_j}(R^\Delta)\,\mathrm{d}\boldsymbol{\Delta}_j = 0$, showing that the slope of the bonding curve, $\mathrm{d}\boldsymbol{\Delta}_j/\,\mathrm{d}\boldsymbol{\Delta}_i = \mathrm{MRS}_{ij}(R^\Delta)$.

tokens that constitutes a share $\alpha - 1$ of their pre-emission supply. The initial emission of LP tokens is chosen arbitrarily by most CFMMs.

The CFMM also maintains the composition of the pool when reserves are withdrawn. This case of negative liquidity provision can be represented as before but with $\alpha \in [0, 1]$. By burning a share $1 - \alpha$ of liquidity tokens in circulation, a liquidity provider receives $O = R(1 - \alpha)$ assets back from the CFMM, and moves its reserves from R to $R' = \alpha R \leq R$.

2.3 Analogy with consumer theory

In consumer theory, the impact of a price change divides into a substitution effect that moves consumption along the consumer's indifference curve, and an income effect that shifts the indifference curve. By contrast, with CFMMs, the substitution and income channels are split between two types of market participants: Traders move reserves along the bonding curve of the CFMM, while liquidity providers shift the bonding curve (see Figure 1). We unfold the implications of this decomposition in the rest of paper, showing that it sheds a powerful light on the inner workings of CFMMs.

Some CFMMs implement more sophisticated operations than simple trade or liquidity provision and withdrawal. Yet, these can be captured by a composition of the two benchmark operations. For example, trading under transaction fees can be seen as a trade followed by a liquidity provision since collected fees are directly fed into the liquidity pool, leading to an upward shift of the bonding curve. Also, non-proportional or partial (e.g. single-sided) liquidity provision triggers a proportional glide along the bonding curve followed by a shift.

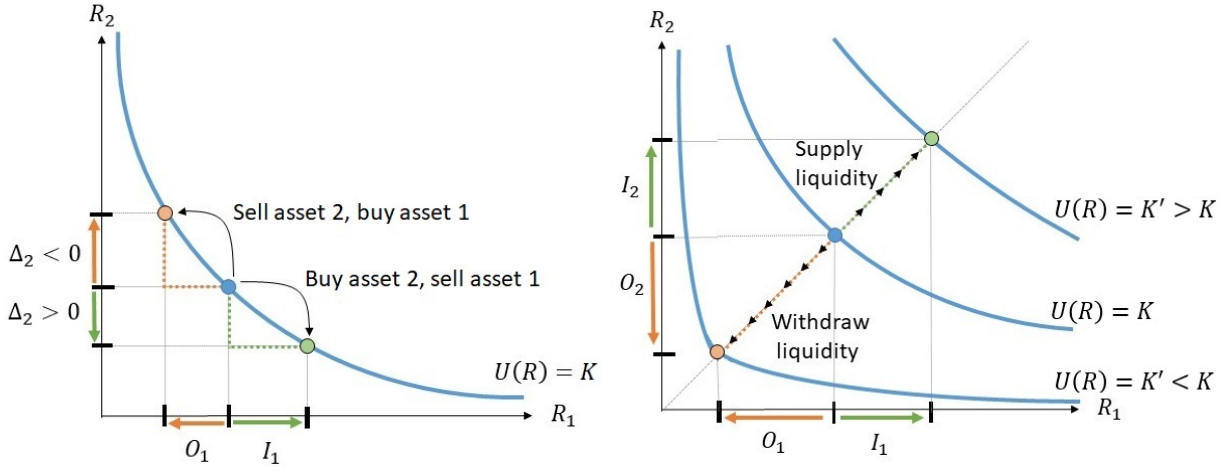

(a) Swap (b) Liquidity provision

Figure 1: Division of CFMM operation among market actors

2.4 Overview of popular CFMMs

Before diving into the mathematical analysis of CFMMs, we survey the trading functions implemented by the most popular protocols (see also Xu et al. 2023 for an extensive survey of existing CFMMs). We restrict our attention to CFMMs that quote prices only as a function of their reserves. Thus we do not cover recent CFMMs that feed additional data to their pricing algorithm by implementing on-chain oracles, like DODO, or that implement concentrated liquidity positions, like Uniswap V3.

Uniswap (Adams et al., 2020)

Uniswap (V2) is by far the most popular CFMM. It was originally deployed on the Ethereum blockchain. Since then, Uniswap has been replicated by multiple clones on Ethereum as well as on other blockchains (e.g., Sushiswap on Ethereum, Pancakeswap on BNB Smart Chain, Serum on Solana). Uniswap's pools contain pre-defined asset pairs and implement the *constant-product* trading function, $U : \mathbb{R}^2_+ \to \mathbb{R}$, such that

$$U(R) = R_1 R_2. \tag{8}$$

Hence a trade is admissible according to Eq. (2) if and only if $(R_1 - \boldsymbol{\Delta}_1)(R_2 - \boldsymbol{\Delta}_2) \geq R_1 R_2$.

Uniswap's LP tokens are called Uniswap (UNI) tokens. Uniswap allows liquidity providers to supply liquidity arbitrarily but advises them to perform proportional multi-asset liquidity provision to avoid arbitrage (see Section 2.2).[8] The initial emission of UNI tokens is set at the geometric mean, $\sqrt{I_1 I_2}$, of the initial input of reserves.

Balancer (Martinelli and Mushegian, 2019)

Balancer is a multi-asset generalization of Uniswap. It uses a *geometric mean* (G3M) trading function, $U : \mathbb{R}^N_+ \to \mathbb{R}_+$, such that

$$U(R) = \prod_{i=1}^N R_i^{w_i}; \qquad w_i \in (0,1), \quad \sum_{i=1}^N w_i = 1. \tag{9}$$

The trading functions of Uniswap and Balancer are both instances of Cobb-Douglas utility functions. Liquidity provision in Balancer can be both proportional multi-asset provision and single-asset provision. In single-asset provision, providers are free to supply any asset individually but pay a trading fee as the CFMM treats this operation as a multi-asset provision followed by a swap.

[8] https://docs.uniswap.org/contracts/v2/guides/smart-contract-integration/providing-liquidity.

MStable (Andersson, 2020)

The MStable CFMM uses a linear trading function, $U : \mathbb{R}_+^N \to \mathbb{R}_+$, with

$$U(R) = \sum_{i=1}^{N} R_i. \tag{10}$$

This *constant sum* CFMM executes swaps at parity. Hence it is convenient for stablecoin pools, since stablecoins pegged to the same underlying value are supposed to trade close to parity. The linear CFMM is not practical because, due to its linearity, arbitrageurs can profit from draining reserves of some asset to zero, thereby preventing other liquidity traders to buy them (see Section 4.2). However, linear trading functions can have a meaningful use when mixed with other functional forms, as in the next CFMM.

Curve (Egorov, 2019a,b)

Curve (previously known as StableSwap) mixes a linear and a geometric-mean trading function.[9] The CFMM's utility is set to simultaneously satisfy $\sum_i^N R_i = U$ and $\prod_i^N R_i = (U/N)^N$. Multiplying the linear trading function by $\chi(U,R)U^{N-1}$ and adding both sides of each identity gives the implicit equation of Curve's trading function, $U : \mathbb{R}_+^N \to \mathbb{R}_+$ such that

$$\chi(R,U)\ U^{N-1}\ \sum_{i=1}^{N} R_i + \prod_{i=1}^{N} R_i = \chi(R,U)\ U^N + \left(\frac{U}{N}\right)^N. \tag{11}$$

The leverage $\chi(U,R)$ is a 0-homogeneous (scale free) function that gives the weight on the linear component. It turns Curve into a geometric CFMM for $\chi(U,R) = 0$ and into a linear one for $\chi(U,R) \to \infty$. The leverage is multiplied by U^{N-1} so that the linear and the geometric component are measured in the same units of utility. The implicit bonding curve originating from Eq. (11) is an hyperbolic-like function, similar to the ones used by Uniswap and Balancer, with a flat central region and asymptotes along each R-axis.

Curve adjusts the leverage dynamically as a function of the reserves:[10]

$$\begin{aligned} &\chi(R,U) = A\chi_0(R,U), \quad \chi_0(R,U) \equiv \frac{\prod_{i=1}^{N} R_i}{(U/N)^N} && \text{in Curve V1;} \\ &\chi(R,U) = A\chi_0(R,U)\frac{\gamma^2}{\gamma + 1 - K_0} && \text{in Curve V2.} \end{aligned} \tag{12}$$

The parameter $A \in \mathbb{R}_+$ is called the amplification coefficient. It controls the flatness of the

[9] Port and Tiruviluamala present a general technique to mix linear and geometric trading functions.

[10] An alternative formulation for Curve V1 and V2's implicit equation follows by plugging the value of $\chi(R,U)$ into Eq. (11): $AN^N \sum_{i=1}^{N} R_i + U = UAN^N + U^{N+1}/N^N \prod_{i=1}^{N} R_i$.

bonding curve around the 45-degree line. The larger is A, the wider is the central plateau of the curve. $\gamma \in \mathbb{R}_+$ controls the steepness of tails by stretching the bonding curve towards the axes when increased (see Fig. 2).

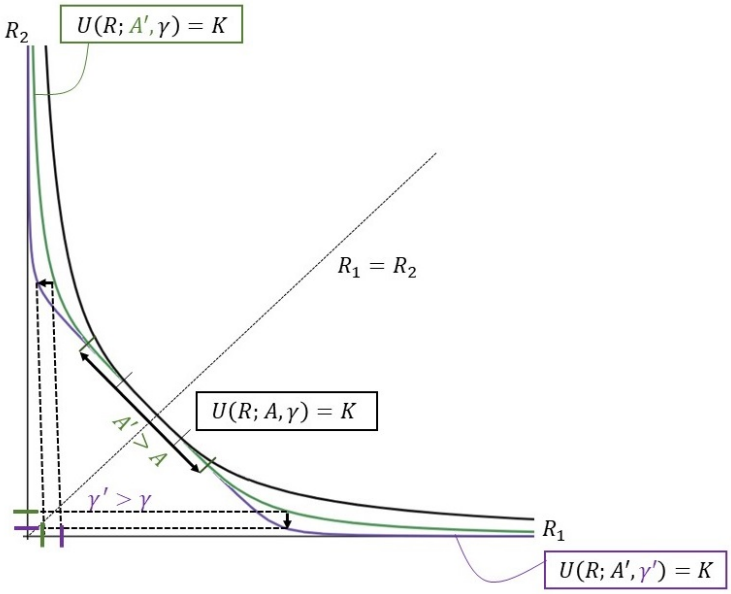

Figure 2: Effect of parameters on Curve's bonding curves

Yieldspace (Niemerg et al., 2020)

Yieldspace uses a *constant power sum* trading function, $U : \mathbb{R}_+^2 \to \mathbb{R}_+$, with

$$U(R) = R_1^\rho + R_2^\rho, \quad \rho \in (-\infty, 1]. \tag{13}$$

This trading function converges to constant product for $\rho \to 0$, and becomes linear at $\rho = 1$. For $\rho \to -\infty$, $U(R)$ converges to a Leontief trading function $U(R) = \min(R_1, R_2)$. In this limit, the CFMM has the bizarre behavior of allowing only trades that switch the reserves of the two assets.[11] That is, assuming $R_1 = \min(R_1, R_2)$, only $\boldsymbol{\Delta} = (R_2 - R_1) \cdot (1, -1)$ is an admissible trade. An immediate multi-asset generalization of the constant power sum is the constant elasticity of substitution (CES) trading function

$$U(R) = \left(\sum_{i=1}^{N} R_i^\rho \right)^{\frac{1}{\rho}}, \quad \rho \in (-\infty, 1). \tag{14}$$

[11] In our notation, the symbol "$\cdot$" denotes the inner product among two vectors. Conversely, two variables next to each other indicate multiplication by a scalar or matrix product.

To the best of our knowledge, this specification has not yet been implemented for $N > 2$.

3 No-Arbitrage Equilibrium

The first question raised by decentralized exchanges is whether they enable discovery of fundamental prices. Given that CFMMs coexist with centralized exchanges which update their quotes at a much higher frequency, it is reasonable to assume that traders observe reference prices $P \in \mathbb{R}^N_+$ for the CFMM's assets. Arbitrageurs can therefore maximize their profits from trading between the CFMM and the centralized market by solving

$$\max_{\boldsymbol{\Delta}} P \cdot \boldsymbol{\Delta} \qquad s.t.\ U\left(R - \boldsymbol{\Delta}\right) \geq K. \tag{15}$$

Problem (15) is not yet in a standard form because it depends on the net trade $\boldsymbol{\Delta}$. However, observing that the arbitrageur's profit is given by

$$P \cdot \boldsymbol{\Delta} = P \cdot \left(R - R'\right) = P \cdot R - P \cdot R', \tag{16}$$

and that $P \cdot R$ is fixed, we infer that the arbitrage problem is equivalent to

$$\min_{R'} P \cdot R' \qquad s.t.\ U\left(R'\right) \geq K. \tag{AP}$$

Two insights follow from formulation (AP) of the arbitrage problem. First, arbitrage profits are maximized when the post-arbitrage value of the pool is minimized. Hence, the incentives of arbitrageurs are opposite to those of liquidity providers. Second, the arbitrageur's problem is isomorphic to the expenditure minimization problem (EMP) of consumers. Thanks to this similarity, an arsenal of established results from consumer theory readily characterizes the solution to the arbitrage problem. Most importantly, it follows that the reserves that solve (AP) are given by the Hicksian demand of the corresponding consumer problem, and liquidity providers' portfolio value of holding the equilibrium reserves in the liquidity pool is described by the related expenditure function. For this reason we refer to $h(P, K)$ as the vector of *Hicksian reserves*, and denote their *portfolio value* by $V(P, K) = P \cdot h(P, K)$. These two quantities are formally defined as

$$h(P, K) = \arg\min_{R} \left\{ P \cdot R \middle| U\left(R\right) \geq K \right\}, \tag{17}$$

$$V(P, K) = \min_{R} \left\{ P \cdot R \middle| U\left(R\right) \geq K \right\} = P \cdot h(P, K). \tag{18}$$

We will refresh these concepts and show exactly how they relate to (AP) in the next para-

graphs, dropping the notation K when irrelevant or clear from the context.

3.1 CFMMs as decentralized price oracles

We are now in a position to derive the prices that should be quoted by the CFMM to reveal the reference values of its assets. When price revelation occurs, the CFMM is said to act as a *decentralized price oracle*. Efficient quotation follows from another well known result in consumer theory according to which, at the solution $R = h(P, K)$ of (AP), the gradient ∇U is proportional to the price vector. That is,

$$P = \lambda \nabla U\left(R\right), \tag{19}$$

where $U(R) = K$ and $\lambda \in \mathbb{R}_+$ is a scaling constant that depends on the choice of numéraire. Assumption 1 ensures that $h(P, K)$ is well-defined, unique, and continuous in P and K. Moreover, condition (19) is both necessary and sufficient for a constrained minimum.

There are two ways to think about Eq. (19). First, it defines how relative prices should be quoted by the CFMM to prevent arbitrage, namely as ratios of marginal utilities as in Eq. (7). Second, it guarantees that arbitrageurs bring each CFMM spot price back in line with the reference price. In other words, arbitrageurs synchronize on-chain and off-chain data, making the CFMM a reliable price oracle.

The decentralized oracle property is the combination of the two above considerations. Using the handy notation $p_{ij}(R) \equiv p_{ij}(0, R)$ to denote CFMM spot prices, it reads:[12]

$$p_{ij}(R) = \frac{U_{R_i}(R)}{U_{R_j}(R)} = \frac{P_i}{P_j}. \tag{20}$$

3.2 Basic properties of Hicksian reserves and portfolio value

Consumer theory connects Hicksian reserves to their shares in the portfolio value. It is well known that the Hicksian reserves of the Cobb-Douglas trading function in Eq. (9) satisfy

$$P_i h_i(P, K) = w_i V(P, K). \tag{21}$$

The exponents w_i are therefore the share of portfolio value that Balancer's liquidity providers obtain from the reserves of asset i. A similar relationship establishes that, for the CES utility

[12] Here $p_{ij}(R)$ is the CFMM spot price, while $p_{ij}(0, R)$ is the price function in Eq. (7) evaluated at $\Delta = 0$. They are different objects: $p_{ij}(R)$ gives the CFMM price as a function of equilibrium reserves; $p_{ij}(\Delta, R)$ gives the updated CFMM price as a function of the purchase Δ of asset i, *given* reserves R.

in Eq. (14),

$$P_i h_i(P,K) = \frac{P_i^{\rho/(\rho-1)}}{\sum_{j=1}^N P_j^{\rho/(\rho-1)}} V(P,K).$$

Hence asset i's share of the portfolio value is now an increasing function of its price. The portfolio value of a linear CFMM comes instead only from the cheapest asset, so that $V(P,K) = K\min(P)$. This is the result of arbitrageurs fully draining the reserves of the most valuable asset.

More generally, for CFMMs that satisfy Assumption 1, standard results (see, for instance, Mas-Colell et al. 1995) ensure that:

Proposition 1. *The Hicksian reserves $h(P,K)$ is: (i) 0-homogeneous in P; (ii) increasing in K ($D_K h(P,K) \geq 0$).*

Proposition 2. *The portfolio value $V(P,K)$ satisfies: (i) 1-homogeneity in P; (ii) $D_P V(P,K) \geq 0$; (iii) $D_K V(P,K) > 0$; (iv) continuity in P and K.*

The homogeneity translates into $h(\cdot)$ being a function of relative, rather than absolute, prices and V scaling linearly with the price vector. The signs of the derivatives above are all very natural: higher utility requires more reserves, and so portfolios that correspond to higher utility are worth more. The continuity of V is inherited from the smoothness of U. Notice that we do not describe $D_P h(\cdot)$ at this stage. Doing so requires a deeper analysis of the off-equilibrium consequences of arbitrage, which we perform later in Section 5.

3.3 CFMM equivalence under monotonic transformation

We have seen in Section 2.4 that CFMMs come in all shapes and sizes. However, their diversity is more apparent than real. It can be trimmed down with the help, once again, of consumer theory. One of the fundamental insight of utility representation is that functional forms matter only to the extent that they capture ordinal preferences. Monotonic transformations are therefore irrelevant for utility functions and, by extension, for trading functions. Concretely, consider two CFMMs with trading functions U_1 and U_2. If there exists a monotonic transformation f such that $U_1 = f \circ U_2$, then the two CFMMs are equivalent. For instance, let us compare Uniswap's trading function $U_A = R_1 R_2$ with the trading function of an equally-weighted Balancer pool, $U_B = R_1^{1/2} R_2^{1/2}$. Since $U_B = \sqrt{U_A}$, or $U_B = f \circ U_A$ with $f(x) = \sqrt{x}$, we can conclude that the two CFMMs are equivalent. In Appendix B, we formally show that:

Example 1. *Uniswap and Balancer with $N = 2$ and $w_1 = w_2 = 1/2$ are equivalent CFMMs.*

This equivalence combined with Eq. (21) shows that the portfolio value of Uniswap's liquidity providers is evenly split among the two assets. Going beyond the Uniswap-Balancer analogy, one can show that Balancer is also equivalent to a CFMM that uses the exp-log trading function $U(R) = \exp(\sum_{i=1}^{n} w_i \log(R_i))$, and that N-assets generalizations of Yieldspace (Eq. 13) are equivalent to the CES CFMMs in Eq. (14) for $\rho \geq 0$.

3.4 Price neutral liquidity provision

The design space can be further narrowed down by requiring that liquidity providers do not affect the prices quoted by the AMM. This restriction is natural since it prevents liquidity provision from scrambling the price discovery process. The necessary condition has long been established in consumer theory where it is common to focus on homothetic preferences so as to ensure that consumers with different incomes demand goods in the same proportions, as long as they are facing the same relative prices. Then the income expansion path of demand becomes linear because the slope of the indifference curves stays constant along rays in $\mathbb{R}^N$ originating at $\mathbf{0}$. The implication for CFMMs is that their spot price depends on the composition and not on the size of the pool solely when the trading function is homothetic. Then, and only then, liquidity provision does not create new arbitrage opportunities by tampering with the spot price (see Figure 3).

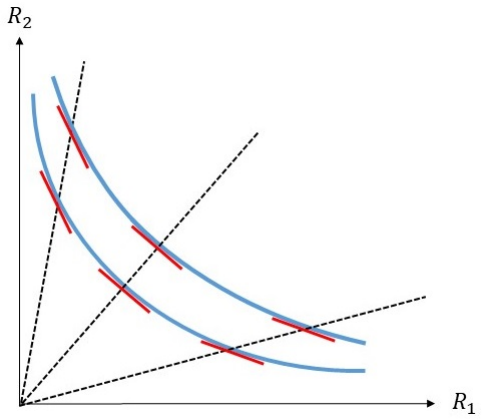

Figure 3: Bonding curves of an homothetic trading function

A trading function is homothetic when it is of the form $U = f \circ u$, where f is an increasing, monotonic function and u is a 1-homogenous trading function. Since we have shown that monotonic transformation are immaterial for the behavior of CFMMs, we can conclude that the class of CFMMs whose trading functions are homogenous of degree one encompasses all the CFMMs whose prices are not impacted by liquidity provision. Imposing this simple restriction therefore results in a dramatic reduction of the design space.

Proposition 3. *CFMM prices are invariant to liquidity provision if and only if U is homothetic. Formally, for $\alpha \in \mathbb{R}$, $D_\alpha p_{ij}(\alpha R) = 0$ if and only if $U = f \circ u$ with f monotone increasing and u 1-homogeneous.*

Proof in Appendix A.

All the CFMMs listed in Section 2.4 are homothetic. Curve is the only CFMM for which homotheticity is not obvious but we show in Appendix B that Curve is indeed 1-homogeneous.

Example 2. *The trading function of Curve is 1-homogeneous.*

4 Traders' Side of the Market

Now that we have established the general structure and properties of CFMMs, we turn our attention to the description of each side of their marketplaces, pinpointing the externalities that connect them. This section explains why the cost for traders of interacting with a CFMM is encapsulated by its price impact. We then show that liquidity providers exert a positive externality on traders when the price impact decreases in the size of the liquidity pool. We provide formal conditions under which this statement is true so that traders really benefit from the participation of liquidity providers.

4.1 Cost of trading and price impact

The cost of buying Δ units of asset i with asset j is given by the difference between the actual transfer and the one that would have prevailed in a market with fixed execution price:

$$C(R^\Delta, R) = q_{ij}(\Delta, R) - p_{ij}(0, R)\Delta, \tag{22}$$

where q_{ij} and p_{ij} are the quantity and price functions defined in Eqs. (6) and (7), while the argument (R^Δ, R) indicates that the reserves moved from R to R^Δ as defined in Eq. (5).[13] According to Eq. (6), the cost of trading is equal to the (total) *price impact*.

$$C(R^\Delta, R) \quad = \quad \int_0^\Delta \Big[p_{ij}(x, R) - p_{ij}(0, R)\Big]\, \mathrm{d}x.$$

We leverage this connection to focus on price impact in order to characterize the cost of trading.

[13] We use (R^Δ, R) instead of (Δ, R) because it underlines the connection between the cost functions of traders and of liquidity providers covered in Section 6.

4.2 Price impact and CFMM curvature

In this subsection, we prove that the CFMMs that satisfy Assumption 1 have strictly positive price impact, so that $D_x p_{ij}(x, R) > 0$. In other words, $q_{ij}(\Delta, R)$ is strictly increasing and convex in Δ for all R.

Our result follows by analogy with the property that a quasi-concave utility function expresses the preferences of an agent with *Diminishing Marginal Rates of Substitution* (DMRS) in consumption, resulting in convex preferences (Arrow and Enthoven, 1961). When we replace the consumer with a CFMM, convex preferences translate into convex trading sets, and DMRS into positive price impact for any swap among *composite* assets generated by linear combinations of the base assets of the liquidity pool.

For simplicity, consider first a two-assets pool. The equivalence between price impact and quasi-concavity in this case follows from decomposing the *curvature* of the bonding curve at reserves R^Δ, given by $s_{ij}(R^\Delta) \equiv D_\Delta p_{ij}(\Delta, R)$, into

$$s_{ij}(R^\Delta) = -\left(\frac{U_{R_iR_i}U_{R_j}^2 - 2U_{R_iR_j}U_{R_i}U_{R_j} + U_{R_jR_j}U_{R_i}^2}{\left(U_{R_j}\right)^3}\right)\Bigg|_{R=R^\Delta} > 0, \tag{23}$$

The expression in parenthesis in Eq. (23) is exactly $D_\Delta \text{MRS}_{ij}(R^\Delta)$. Thus, owing to DMRS, $s_{ij}(R^\Delta) > 0$ for all Δ and R.

Lemma 1. *A two-asset CFMM exhibits positive price impact if and only if the trading function is strictly quasi-concave.*

Proof in Appendix A.

We provide additional details on the equivalence between DMRS and quasi-concavity in Appendix D.1. For a generic N-asset CFMM, quasi-concavity becomes apparent when we define a trading function over pairs of linearly-independent (i.e. non proportional) composite assets $\mathbf{b}, \mathbf{b}' \in \mathbb{R}^N$ with reserves $R_\mathbf{b}, R_{\mathbf{b}'} \in \mathbb{R}$. Linear-independence is needed to consider the two composite assets as distinct, so that $\mathbf{b}$ and $\mathbf{b}'$ form an orthogonal basis for the swap. The trading function so constructed, $\widetilde{U} : \mathbb{R}^2 \times \mathbb{R}^N \times \mathbb{R}^N \to \mathbb{R}$, is defined as

$$\widetilde{U}(R_\mathbf{b}, R_{\mathbf{b}'}; \mathbf{b}, \mathbf{b}') = U(R_\mathbf{b}\mathbf{b} + R_{\mathbf{b}'}\mathbf{b}'), \text{ with } \mathbf{b} \neq \alpha\mathbf{b}', \text{ and } \alpha \in \mathbb{R} \setminus \{0\}. \tag{24}$$

Using $\widetilde{U}$ we can generalize the concept of quasi-concavity to more than two assets since Arrow and Enthoven (1961) have shown that U is quasi-concave if and only if $\widetilde{U}$ is quasi-concave for *all* (linearly-independent) assets $(\mathbf{b}, \mathbf{b}') \in \mathbb{R}_+^N \times \mathbb{R}^N$.

Proposition 4. *A CFMM exhibits positive price impact among composite assets if and only if the trading function is strictly quasi-concave.*

An alternative and more common way to test for quasi-concavity of a function with $N \geq 3$ arguments is by studying the sign of the principal minors of its bordered Hessian.[14] The two approaches are equivalent, although the method we propose has a direct economic interpretation when applied to CFMMs: Identifying quasi-concavity coincides with identifying positive price impact in the trading of composite assets.

Restricting Proposition 4 to basis vectors $\mathbf{b} = \mathbf{i}$, $\mathbf{b'} = \mathbf{j}$ immediately gives the most relevant application of quasi-concavity in practice:

Corollary 1. *A CFMM exhibits positive price impact among base assets if and only if* $\widetilde{U}$ *in Eq. (24) is quasi-concave for all basis vectors* $\boldsymbol{b} = \boldsymbol{i}$, $\boldsymbol{b'} = \boldsymbol{j}$.

Infinite liquidity

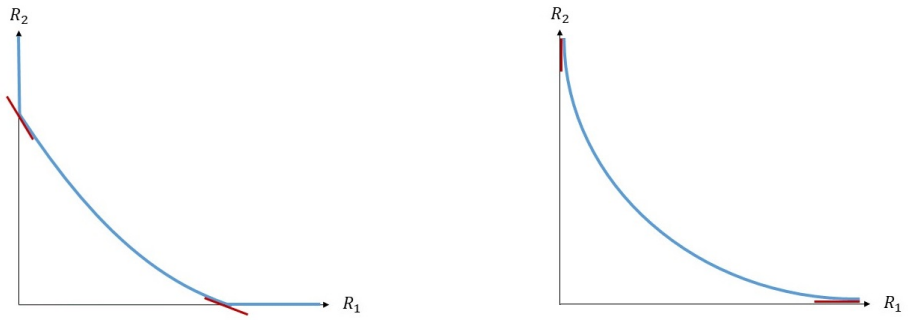

(a) Depletable CFMM (b) Non-depletable CFMM

Figure 4: Infinite liquidity and asymptotes of the bonding curve

A positive price impact is an inconvenience for traders but, as we will see in Section 6.2, it is also a necessary defense mechanism for the CFMM. Besides reducing arbitrageurs' incentives to trade against liquidity providers, it can guarantee the property known as 'infinite liquidity' in practitioners' jargon. That is, the price impact can preclude traders from exhausting all reserves, thereby ensuring that the CFMM always remains ready to process incoming orders. For this to be the case, strict quasi-concavity has to be complemented with the property that the price becomes infinite when the reserve of any of the traded assets nears zero:

[14] This is also known as the determinant criterion: U is quasi-concave if and only if the $(-1)^r |\bar{D}_r^2 U| \geq 0$ for $r = 1, 2, \ldots, N$, where $\bar{D}_r^2 U$ is the r-th principal minor of the bordered Hessian of U.

Lemma 2. *The CFMM's reserves cannot be drained if*

$$\lim_{\Delta_i \to R_i} \left| MRS_{ij}(R^{\Delta}) \right| = +\infty, \quad \textit{for all } ij. \tag{25}$$

Eq. (25) is a variant of the *Inada condition*. It essentially requires that the bonding curve has asymptotes along the axis of each asset pair, as in Fig. 4b. By contrast, bonding curves such as those in Fig. 4a correspond to CFMMs that can be completely drained of their reserves.

4.3 Network effect of liquidity provision on trade

The results derived in the previous subsection allow us to characterize how the cost of trading $C(R^{\Delta}, R)$ is affected by liquidity provision. We focus on multi-asset proportional liquidity provisions to homothetic CFMMs since they are by far the most common in practice (see Sections 2.2 and 3.4). We relegate the more convoluted analysis of single-asset liquidity provision to Appendix C.

Equilibrium spot prices are preserved under multi-asset provision to an homothetic CFMM. Thus, liquidity provision has a positive network effect on traders whenever purchasing assets at the new reserves costs less inputs than before. Specifically, traders benefit from a proportional liquidity injection if

$$q_{ij}(\Delta, R) - q_{ij}(\Delta, \alpha R) \geq 0, \quad \alpha > 1, \tag{26}$$

holds for all distinct asset-pairs ij and strictly for at least one ij. Fig. 5 shows an example where Eq. (26) is satisfied.

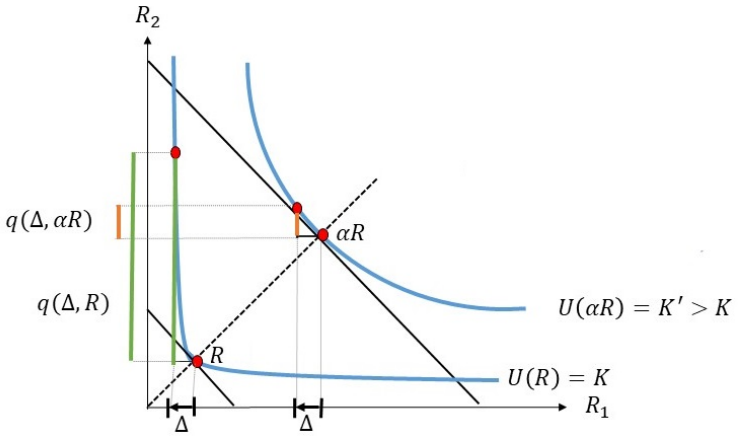

Figure 5: Effect of deepening the pool on the cost of trading

In Appendix A we establish a property which allows us to conclude that liquidity provision does exert a positive externality on traders for homothetic trading functions under Assumption 1. Namely, we show that homotheticity of the trading function implies 1-homogeneity of the quantity function, and vice-versa:

Lemma 3. *A trading function $U(R)$ is homothetic if and only if its quantity function $q_{ij}(\Delta, R)$ is 1-homogeneous in (Δ, R).*

Furthermore, it follows from Euler's theorem that:

Corollary 2. *If U if homothetic, then the price function $p_{ij}(\Delta, R)$ is 0-homogeneous in (Δ, R) so that*

$$p_{ij}(\Delta, \alpha R) = p_{ij}\left(\frac{\Delta}{\alpha}, R\right).$$

In words, buying Δ in a pool that is deeper by a factor α is equivalent to buying the smaller quantity Δ/α at the initial reserves. Thus, under quasi-concavity of the trading function, $p_{ij}(\Delta, \alpha R) < p_{ij}(\Delta, R)$ due to Corollary 1, and so, $q_{ij}(\Delta, \alpha R)$ being determined through integration of the marginal price, we have $q_{ij}(\Delta, \alpha R) < q_{ij}(\Delta, R)$.

Proposition 5. *If the trading function U is homothetic and strictly quasi-concave, liquidity provision reduces the cost of trading; i.e. the inequalities in Eq. (26) hold strictly for all asset pairs.*

Proof in Appendix A.

5 Liquidity Providers' Side of the Market

While the previous section focused on the externalities affecting traders, we now turn our attention to liquidity providers. In centralized markets managed via limit order books, liquidity providers bear an adverse selection cost (Glosten and Milgrom, 1985). CFMMs are no exception. We have shown in Section 3 that arbitrage opportunities harm liquidity providers because arbitrageurs replace appreciating assets with depreciating ones. To quantify the impact of arbitrage, new analytical tools are needed. It turns out that the effect of arbitrage on the portfolio value of liquidity providers is more conveniently studied in the *dual space* of the arbitrage problem. The benefit of this shift of perspective is intuitive since the preferences of liquidity providers are not expressed by the trading function but rather by the monetary value encoded in its supporting hyperplanes. Moreover, formulating our analysis in the dual

space does not entail any loss of information because the Hicksian reserves can be derived from the portfolio value function through Shephard's Lemma:[15]

Proposition 6. *(Shephard's Lemma) Suppose that $U : \mathbb{R}^N_+ \to \mathbb{R}$ is a continuous, strictly increasing, and strictly quasi-concave trading function. Then, for all $P \in \mathbb{R}^n_{++}$ and $K \in \mathbb{R}$,*

$$h\left(P, K\right) = \nabla_P V\left(P, K\right), \tag{27}$$

where $h\left(P, K\right)$ and $V\left(P, K\right)$ are the Hicksian reserves and their portfolio value, defined in Eqs. (17) and (18).

Furthermore, according to Eq. (27), twice-differentiation of the portfolio value with respect to P yields the impact of price movements on equilibrium reserves. The resulting Hessian is the *Slutsky matrix* which encapsulates the substitution effects of a change in prices on consumption.

Corollary 3. *The price derivative of the Hicksian reserves are given by the Slutsky matrix: $D^2_P V(\cdot, K) = D_P h(\cdot, K)$.*

We refer to the negative of the ij-element of the Slutsky matrix

$$\ell_{ij}(P, K) \equiv -D^2_P V_{ij}(\cdot, K) = -\partial h_i(P, K)/\partial P_j \tag{28}$$

as the *liquidity* of asset i with respect to the price of asset j. Liquidity is a standard concept to express how many units of R_i the CFMM can sell in response to a marginal change in P_j (Goyal et al., 2022; Milionis et al., 2022). This interpretation holds for $\ell_{ij}(P, K) \geq 0$. Similarly, in the case of negative liquidity, with $\ell_{ij}(P, K) \leq 0, \left|\ell_{ij}(P, K)\right|$ gives the amount of R_i that the CFMM can buy in response to a marginal change in P_j. Large values of $\ell_{ij}(P, K)$ also indicate that the CFMM can sell (or buy) many reserves without altering substantially its internal spot prices, thereby remaining synchronized with the reference market prices. Besides describing how reserves change with arbitrage, we will see in Section 6.2 that liquidity connects the perspectives of liquidity providers and traders.

5.1 Liquidity and concavity of portfolio value

The concavity of the portfolio value quantifies the exposure to adverse selection of liquidity providers. Technically, $V(P, K)$ is concave if, given a pair of distinct prices, the portfolio

[15] In consumer theory, Shephard's Lemma is typically stated assuming local non-satiation rather than strict monotonicity of the utility function. The latter is a stronger assumption.

value evaluated at the average price is higher than the average portfolio value at each separate price:

$$V(\alpha P + (1-\alpha)P', K) \geq \alpha V(P, K) + (1-\alpha)V(P', K), \text{ for all } P, P', \text{and } \alpha \in [0,1]. \quad (29)$$

We can interpret Eq. (29) as comparing the portfolio value at the average prices to the expected value of the portfolio resulting from a distribution that returns P with probability α and P' with probability $1-\alpha$. Intuitively, the inequality in Eq. (29) holds because the arbitrageur can re-optimize reserves at each price level, generating a profit for herself and a corresponding loss for the liquidity providers. Hence the expected portfolio value over multiple prices should be lower than the portfolio value at the fixed average price. Assumption 1 guarantees that arbitrage is indeed profitable after each price change.

Lemma 4. *The portfolio value $V(P, K)$ is concave in P.*

Proof in Appendix A.

Liquidity

The combination of Lemma 4, Proposition 6 and Corollary 3 allow us to translate the concavity of $V(P, K)$ into conditions on the signs and magnitudes of assets' liquidity:

Proposition 7. *$D_P^2 V(\cdot, K)$ is negative semi-definite and satisfies $D_P^2 V(\cdot, K)P = \mathbf{0}$.*

Proposition 7 essentially implies that prices and reserves move in opposite directions, a behavior encapsulated in consumer theory by the *compensated law of demand*.[16] In particular, semi-definite $D_P^2 V(\cdot, K)$ and $D_P h(\cdot, K)P = \mathbf{0}$ jointly imply that

$$\ell_{ii}(P, K) \geq 0 \text{ for all assets } i, \text{ and } \ell_{ij}(P, K) \leq 0 \quad \text{for at least one } j \neq i. \quad (30)$$

In words, an increase in P_i that alters the no-arbitrage equilibrium allows for liquidity extraction from the reserves of asset i but requires a liquidity injection of at least another asset j to maintain reserves on the bonding curve.

5.2 Impermanent losses

We are now in a position to quantify the cost of adverse selection originating from arbitrage. Our measure of adverse selection is the profit that the arbitrageur can make from a single

[16] The law of demand states that marginal changes in prices $\mathrm{d}P \in \mathbb{R}^N$ and reserves $\mathrm{d}h(P, K) \in \mathbb{R}^{2N}$ satisfy $\mathrm{d}P \cdot \mathrm{d}h(P, K) = \mathrm{d}P \cdot D_P h(\cdot, K)\,\mathrm{d}P \leq 0$.

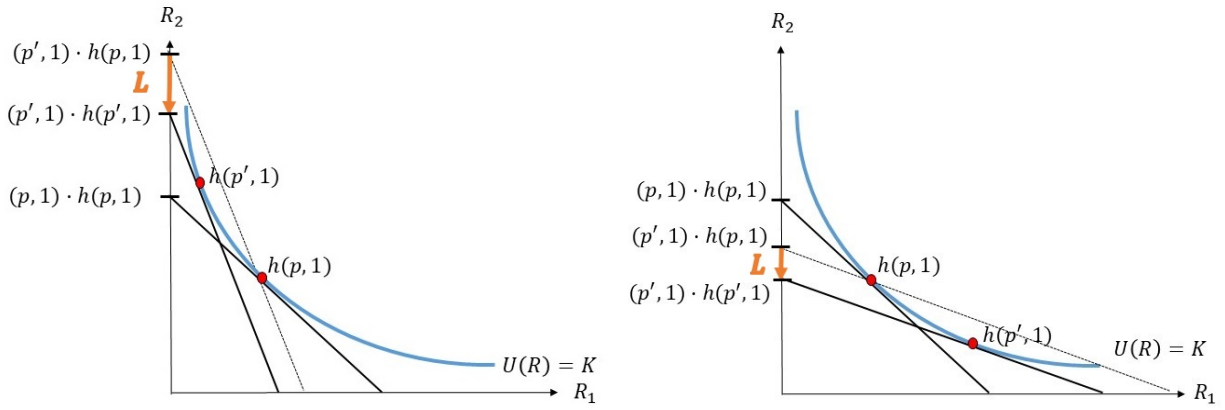

(a) $p' > p$ (b) $p' < p$

Figure 6: Impermanent losses in the primal space

price movement. The cost component originating from arbitrage is known as *impermanent loss*.[17] To quantify it, consider a change in reference market prices from P to P', assuming that the spot prices of the CFMM were initially equal to their reference value P. Before arbitrageurs consume their arbitrage opportunity, the value of the liquidity pool is $P' \cdot h(P,K)$. However, at the new equilibrium, the pool's resources will be worth $V(P',K) = P' \cdot h(P',K) \leq P' \cdot h(P,K)$. The liquidity providers therefore lose

$$L(P',P,K) = P' \cdot \Big(h(P,K) - h(P',K)\Big) \geq 0. \tag{31}$$

The inequality follows directly from the definition of the portfolio value since it achieves a global minimum at any given prices. Moreover, Eq. (31) holds with equality at $P' = \alpha P$ (for $\alpha > 0$) as $h(\alpha P, K) = h(P,K)$.

Fig. 6 represents the impermanent loss in the primal space of the arbitrage problem (AP). It depicts the loss generated by a change from p to p' in the price of asset 1 measured in terms of the numéraire asset 2. With this normalization, the price vector is $P = (p,1)$, with $p = P_1/P_2 = P_1$. The impermanent loss L reported in orange on the vertical axis is the difference between the ordinate-intercepts of the hyperplanes with slope p' that are tangent to the bonding curve at $h(p,1)$ and $h(p',1)$. These correspond to the portfolio values of liquidity providers before and after arbitrage. We now show that representing the impermanent loss in the dual space provides a more compact and transparent representation.

[17] An alternative term for the adverse selection cost of liquidity provision is divergence loss (Goyal et al., 2022). While some papers refer to impermanent loss specifically for passive liquidity positions, we adopt it as a generic adverse selection cost as it is more widely used in both the literature and by practitioners.

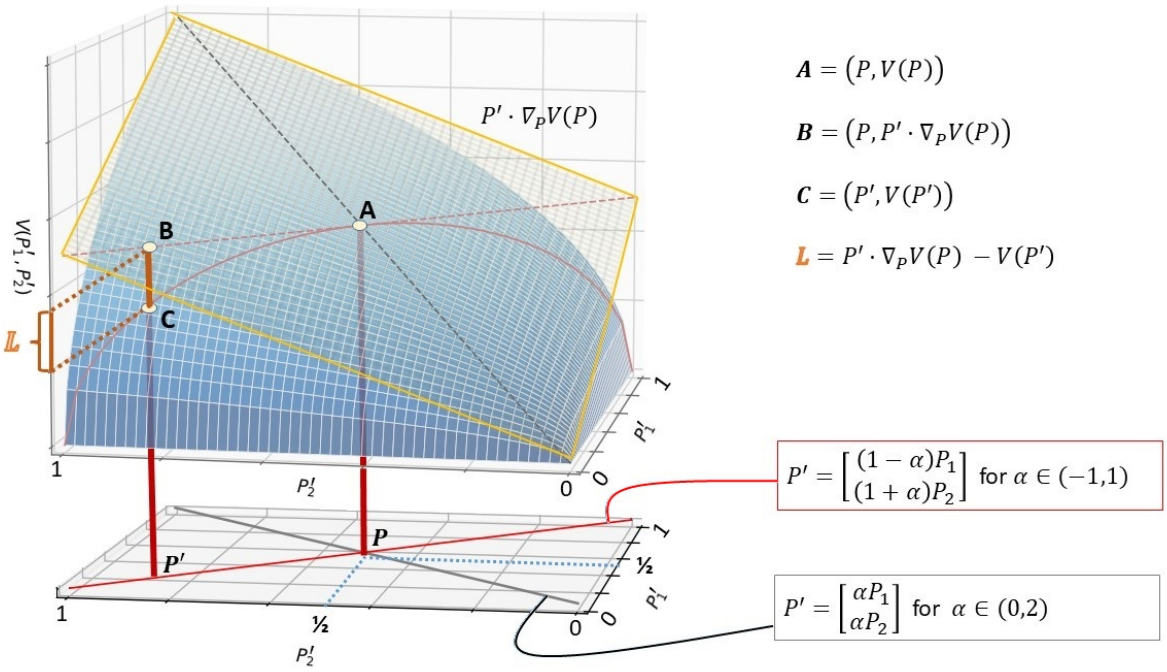

Figure 7: Geometry of impermanent losses in the dual space

Shephard's Lemma (Proposition 6), allows us to rewrite the impermanent loss as

$$L\left(P', P, K\right) = P' \cdot \nabla_P V\left(P, K\right) - V\left(P', K\right). \tag{32}$$

As shown in Fig. 7, $L\left(P', P, K\right)$ is the difference between the portfolio value and its tangent hyperplane at prices P.[18] $L\left(P', P, K\right) \geq 0$ then immediately follows from the concavity of the portfolio value established in Proposition 2. Notice that impermanent losses are strictly positive only when *relative* prices change. Instead, if all prices are scaled by the same factor α, as on the gray price curve in Fig. 7, then the portfolio value moves along the black dashed line on the supporting hyperplane. In other words, the portfolio value also scales linearly from $V(P)$ to $\alpha V(P)$, achieving the same values as those on the supporting hyperplane. By contrast, if relative prices change, as on the red price curve, then the value of $V(P')$ is lower than that of $P' \cdot \nabla_P V(P)$, as can be seen comparing the solid red parabola with the red dashed line. The impermanent loss induced by a price change from P to P' is therefore given by the vertical distance between points **B** and **C** in Fig. 7. Given that impermanent losses are measured by the difference between a linear and a concave function, they are convex in

[18] This representation holds true because of Shephard's Lemma. Omitting the argument K, the difference between the portfolio value $V(P')$ and its approximation based on its supporting hyperplane at P is given by $\left[V(P) + \nabla_P V(P) \cdot (P' - P)\right] - V(P')$, which equals $P' \cdot \nabla_P V(P) - V(P')$ because Shephard's Lemma implies that $V(P) = P \cdot \nabla_P V(P)$.

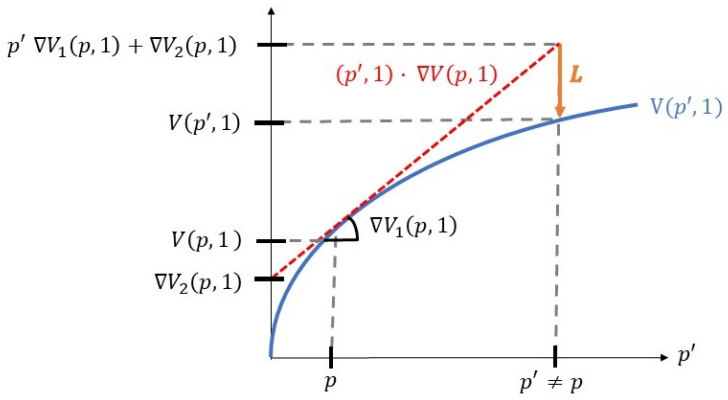

Figure 8: Projected impermanent losses for a change in the relative price

the price displacement. In Appendix A we complement these graphical arguments with an analytical proof which demonstrates that:

Proposition 8. *The impermanent losses $L(P', P, K)$ are positive and convex in P', with global minimum 0 at $P' = \alpha P$, for all $\alpha \in \mathbb{R}_{++}$.*

Fig. 8 reports a 2-dimensional projection of Fig. 7 for the same setting as Fig. 6a. As the figure shows, a change in relative prices causes an impermanent loss for every $p' \neq p$.

Impermanent losses and price-neutral liquidity provision

We have shown in Section 3.4 that the spot prices of an homothetic CFMM are independent of the overall liquidity or utility K of the CFMM. Interestingly, this property also guarantees that the impermanent loss as a percentage of the portfolio value is independent of K. This result holds true because, although liquidity provision incentivizes arbitrage by flattening the curvature of the bonding curve, it also proportionally increases the value of the pool. The proof follows once again from standard microeconomic theory. It is a direct consequence of the fact that homothetic trading functions generate separable portfolio value functions and separable Hicksian reserves.

Lemma 5. *$V(P, K)$ and $h(P, K)$ are separable if and only if U is homothetic; that is, $V(P, K) = \phi(K)\, V(P, 1)$, $h(P, K) = \phi(K) h(P, 1)$, where $\phi(K)$ is a positive and increasing function.*

Proof in Appendix A.

Lemma 5 implies that

$$\frac{L(P', P, K)}{V(P, K)} = \frac{P' \cdot (h(P, K) - h(P', K))}{P \cdot h(P, K)} = \frac{P' \cdot (h(P, 1) - h(P', 1))}{P \cdot h(P, 1)}$$

is independent of K when U is homothetic, as $\phi(K)$ cancels out after being factorized in the numerator and denominator.

Proposition 9. *The rate of impermanent losses, $L(P', P, K)/V(P, K)$, is independent of the pool size K whenever the CFMM is homothetic.*

5.3 Price manipulation

Beyond its theoretical importance, understanding the role of arbitrage in shaping pool values and Hicksian reserves has practical implications. A notable example is the exploit on Warp Finance, which stemmed from a design flaw in its quoting algorithm (Michel, 2022).

Warp Finance allowed borrowing using Uniswap's LP tokens as collateral, and the attack exploited flaws in how Warp Finance calculated the value of these tokens. While the formula for pricing LP tokens can be computed simply as the product of asset prices, P, and reserves, R, divided by the total supply of LP tokens, its implementation requires careful consideration of these inputs are determined. Warp Finance used a time-weighted average price (TWAP) to compute P, effectively mitigating short-term volatility, but relied on instantaneous (current) reserves from Uniswap for R. This inconsistency between the average prices and current reserves created an exploitable vulnerability.

Using a *flash loan*, an attacker manipulated the reserves to inflate the perceived value of LP tokens. A flash loan is a unique DeFi mechanism that allows users to borrow large sums of assets without collateral, provided the loan is repaid within a single composite transaction including multiple operations. If not repaid, the entire transaction is reverted, effectively nullifying the loan. Exploiting this mechanism, the attacker temporarily injected liquidity into the Uniswap pool, increasing its reserves R and artificially inflating the value of the LP tokens. With these overvalued tokens as collateral, the attacker borrowed more funds than the collateral's true worth.

An alternative DeFi protocol, Alpha Finance, addressed this type of vulnerability by incorporating equilibrium-based reserve values, using the Hicksian reserves' formula, into its pricing mechanism. Specifically, rather than relying on current reserves, Alpha Finance relied on an historical average of Hicksian reserves calculated from the constant product invariant of liquidity pools and trusted price oracles for asset prices. This approach anchored LP token valuations to the pool's theoretical equilibrium, making it resilient to flash loans.

The Warp Finance exploit, and Alpha Finance's response, highlight the importance of incorporating equilibrium pricing principles to mitigate vulnerabilities in Constant Function Market Makers (CFMMs) and strengthen the design of resilient DeFi protocols.

6 Duality in Liquidity Provision and Liquidity Trading

We conclude our analysis by showing that duality ties together the perspectives of traders and liquidity providers. First, we outline a method for portfolio replication: Given a portfolio value $V(P, K)$, we explain how to find the trading function U which achieves that value. Then we build on this result to highlight the fundamental tradeoff connecting the two market sides of a CFMM. We prove that the externality exerted by liquidity providers on traders is inversely proportional to the externality exerted by traders on liquidity providers. Moreover, we show that the size of these externalities is fully determined by the slope of the bonding curve.

6.1 Portfolio replication

The construction method for U is equivalent to the procedure for recovering preferences from expenditures. It relies on the dual-conjugacy connection among portfolio value and trading set, as first pointed out in the seminal paper by Angeris et al. (2021a). The construction of the portfolio value starts from the *negative* indicator function of the trading set,

$$\left[-\delta_{\mathcal{S}(K)}\right](R) = \begin{cases} 0 \text{ if } R \in \mathcal{S}(K) \\ -\infty \text{ otherwise} \end{cases}, \tag{33}$$

which differs from the standard indicator function $\delta_{\mathcal{S}(K)}(R)$ in that it returns $-\infty$ instead of ∞ whenever $R \notin \mathcal{S}(K)$. Given $[-\delta_{\mathcal{S}(K)}]$, its Fenchel conjugate $\left[-\delta^*_{\mathcal{S}(K)}\right](P)$ retrieves the portfolio value:[19]

$$\left[-\delta^*_{\mathcal{S}(K)}\right](P) \equiv \inf_R \left(P \cdot R - \left[-\delta_{\mathcal{S}(K)}\right](R)\right) = V(P, K). \tag{34}$$

Notice that duality yields an alternative proof that the portfolio value $\delta^*_{\mathcal{S}(K)}$ is concave in P since it is the pointwise infimum of a family of linear functions.

Likewise, if $\mathcal{S}(K)$ is convex (U is quasi-concave), the Fenchel-Moreau theorem ensures that the trading set can be recovered from the portfolio value via double conjugation. In other words, the biconjugate $[-\delta^{**}_{\mathcal{S}(K)}](R)$ recovers the negative indicator of the trading set:

$$V^*(R, K) \equiv \inf_P \left(P \cdot R - V(P, K)\right) = [-\delta^{**}_{\mathcal{S}(K)}](R) = \left[-\delta_{\mathcal{S}(K)}\right](R).$$

[19] Since $[-\delta_{\mathcal{S}(K)}]$ is a concave function of P, Eq. (34) uses the concave Fenchel conjugate rather than the more common convex conjugate, which is defined differently. For example, the convex conjugate of $\delta_{\mathcal{S}(K)}$ is $\delta^*_{\mathcal{S}(K)} = \sup_P \left(P \cdot R - \delta_{\mathcal{S}(K)}(R)\right)$.

To the best of our knowledge, $\mathcal{S}(K)$ is convex for all the existing CFMMs so recovery via dual conjugacy is widely applicable. If $\mathcal{S}(K)$ is not convex, biconjugation recovers instead its convex closure.

Proposition 10. *The portfolio value and the (negative) indicator of the trading set are dual conjugates:*

$$V(P,K) = \left[-\delta^*_{\mathcal{S}(K)}\right](P), \text{ and } \left[-\delta_{\mathcal{S}(K)}\right](R) = V^*(R,K). \tag{35}$$

Proof in Appendix A.

We give a practical demonstration of this technique in Appendix B where we explain how to recover Uniswap's trading set from its portfolio value function.

Example 3. *Duality recovers Uniwap's CFMM from* $V(P,K) = 2\sqrt{K(P_1P_2)}$.

Angeris et al. (2021a) show that the above replication technique can be used to emulate more sophisticated financial products. For example, they construct a CFMM that replicates a covered call option or a perpetual American put option.

6.2 Tradeoff between liquidity provision and trading

We now use duality to establish an inverse relation between the sensitivity to shocks of impermanent losses and of trading costs. This finding, by proving that the preferences of liquidity providers and traders are divergent, opens the way for a characterization of the Pareto frontier given market fundamentals. For the sake of exposition, we present the tradeoff for a two-asset pool although the analysis can be generalized.

Consider a 2-asset CFMM with initial reserves $R = (R_i, R_j)$ and initial prices $P = (p_{ij}, 1)$, so that asset j is the numéraire and p_{ij} is the relative price of asset i. The expected cost of buying x units of asset i according to a strategy centered around 0 with variance $\mathbb{V}(x)$ satisfies the approximation

$$\mathbb{E}_x\left[C\left(R - x\mathbf{i} + q(x,R)\mathbf{j}, R\right)\right] \approx \frac{1}{2}\,\mathbb{V}(x)\; s_{ij}(R) \tag{36}$$

when x is small relative to the CFMM's reserves. Similarly, the expected impermanent loss induced by a small change y, of mean 0 and variance $\mathbb{V}(y)$, in the price of i satisfies[20]

$$\mathbb{E}_y\left[L\left(P + y\mathbf{i}, P\right)\right] \approx \frac{1}{2}\,\mathbb{V}(y)\; \ell_{ii}(P). \tag{37}$$

[20]Let R^x and P^y denote updated reserves and prices. By Taylor-expanding $C(R^x, R)$ around $x = 0$, we get $\mathbb{E}_x\left[C(R^x,R)\right] = \mathbb{E}(x)\; D_xC(R,R) + \mathbb{E}(x^2)D_x^2C(R,R)/2 + o(\mathbb{E}(x^2))$. The first-order term is zero since $\mathbb{E}(x) = 0$ (and $D_xC(R,R) = 0$). For the second-order term, $\mathbb{E}(x^2)D_x^2C(R,R)/2 = \mathbb{V}(x)s_{ij}(R)/2$. Higher-order terms are negligible if $\mathbb{V}(x)$ is small. The reasoning is analogous for $\mathbb{E}_y\left[L(P^y,P)\right]$.

We now show that, *in equilibrium*, $s_{ij}(R)$ and $\ell_{ii}(P)$ are inversely proportional. Thus, one cannot tune the trading function so as to make traders better off without simultaneously making providers worse off. This is a consequence of the dual nature of equilibrium prices and reserves, which allows us to treat them as inverse functions of each other.

Specifically, in Appendix D.2 we show that, having fixed a numéraire asset, there exists by Assumption 1 a unique price vector given by the *inverse* of the Hicksian reserves $h^{-1} : \mathcal{S}^b(K) \to \mathbb{R}^N_+$, where $\mathcal{S}^b$ is the bonding curve defined in Eq. (3). Since the trading set is convex, the separating hyperplane theorem guarantees that every point on the bonding curve can be uniquely identified by the slope of a supporting hyperplane. Having defined $h^{-1}(R) = (p_{ij}(R), 1)$, we can use the inverse function theorem to establish that the sensitivities $s_{ij}(R)$ and $\ell_{ii}(P)$ in Eqs. (36) and (37) are reciprocal of each other.

Proposition 11. *At the no-arbitrage equilibrium, price impact and liquidity are reciprocal to each other:*

$$\ell_{ii}(P) = \frac{1}{s_{ij}\left(h(P)\right)}, \qquad\qquad s_{ij}(R) = \frac{1}{\ell_{ii}(p_{ij}(R), 1)}.$$

Proof in Appendix A.

Combining Eqs. (36) and (37) with Proposition 11, we see that the cost of trading is proportional to the curvature of the bonding curve, captured by the price impact $s_{ij}(R)$, whereas impermanent losses are proportional to the liquidity $\ell_{ii}(P)$ given by the *reciprocal* of such curvature. Therefore low curvature trading functions incentivize trading while high curvature trading functions incentivize liquidity provision. All CFMMs are subject to this fundamental tradeoff. Hence there cannot be a choice of trading function that dominates others in all respects.

Fig. 9 illustrates this principle by comparing the impact of a price shock across two CFMMs with different curvatures. Both panels show the impermanent losses resulting from the response of the arbitrageur. The left-hand side panel uses a convex bonding curve while the right-hand side panel uses a linear bonding curve. Comparing $L(P', P)$ and $q(\Delta, R)$ in the two panels, one can see that more curvature reduces the impermanent loss for liquidity providers but increases the cost of trading.

7 Transaction Fees

In this section, we show how to incorporate transaction fees in the prior analysis. As Angeris et al. (2021b) shows for Uniswap, a transaction fee $\tau \in (0, 1)$ on asset exchanges can be incorporated into the CFMM by modifying condition (4) to scale the inputs by a factor of

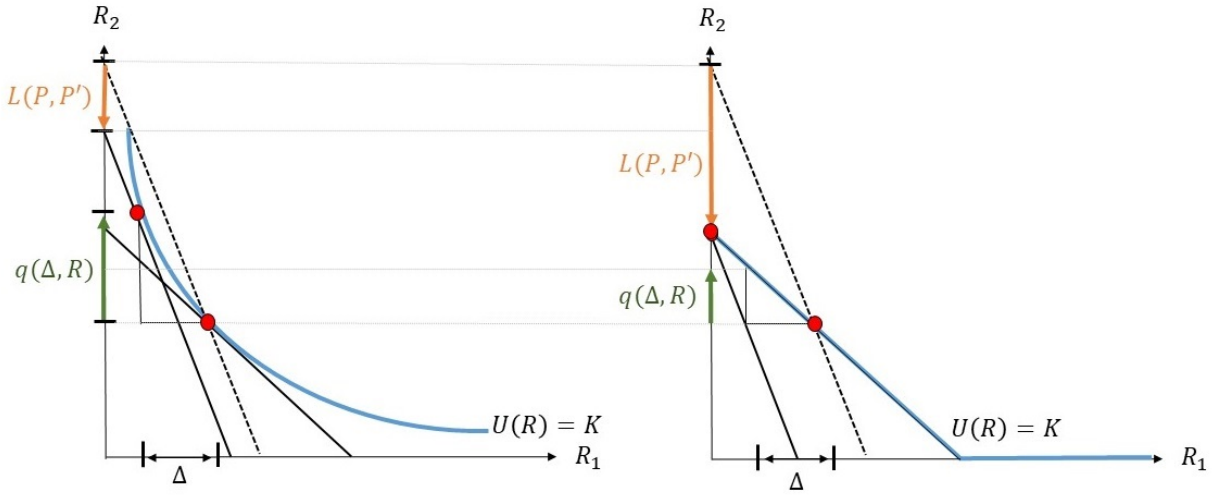

Figure 9: Tradeoff between impermanent losses and trading costs

$(1-\tau)$. In this way, a fraction τ of the inputs only accrues to the liquidity pool without counting as input for the asset exchange. The condition for withdrawing Δ_i units of asset i in exchange for Δ_j units of asset j becomes

$$U(R - \boldsymbol{\Delta}_i + (1-\tau)\boldsymbol{\Delta}_j) = U(R), \tag{38}$$

where $\boldsymbol{\Delta}_i$ is a vector representing the change in reserves of asset i (which is non-zero only for asset i) and $\boldsymbol{\Delta_j}$ is defined identically. Condition (38) can be used to determine the cost of buying asset i in terms of asset j units as well as the revenue from selling asset j in terms of asset i units.

By implicitly differentiating Eq. (38) as in the derivation of Eq. (7), one can see that transaction fees create a bid-ask spread. The bid (buying) price for asset i is given by:

$$p_{ij}^{\text{bid}}(R;\tau) = p_{ij}(R) \times \frac{1}{1-\tau},$$

where $p_{ij}(R)$ is the spot price of asset i quoted by a CFMM without fees, defined in Eq. (20). Conversely, the ask (selling) price of asset i is given by:

$$p_{ij}^{\text{ask}}(R;\tau) = p_{ij}(R) \times (1-\tau).$$

The resulting bid-ask spread is thus

$$p_{ij}^{\text{bid}}(R;\tau) - p_{ij}^{\text{ask}}(R;\tau) = p_{ij}(R) \times \left[\frac{1}{1-\tau} - (1-\tau)\right] \approx p_{ij}(R) \times \tau$$

for τ sufficiently small, as is typically the case in practice.[21]

No-Arbitrage Equilibrium with transaction fees

The addition of transaction fees causes the CFMM to track the reference market price only within a no-arbitrage region. When the reference market price lies within the bid-ask spread, no trade with the CFMM can be profitable. More precisely, let m_{ij} denote the market price of asset i in terms of asset j. An arbitrageur can make a profit by buying asset i from the CFMM and selling it to the reference market only if $m_{ij} > p_{ij}^{\text{bid}}(R;\tau)$. On the other hand, the arbitrageur can profit from buying asset i from the reference market and selling it to the CFMM only if $m_{ij} < p_{ij}^{\text{ask}}(R;\tau)$. So the equilibrium no-arbitrage region is such that

$$\frac{m_{ij}}{p_{ij}(R)} \in \left[1-\tau, \frac{1}{1-\tau}\right].$$

8 Profits and Losses from Liquidity Provision

The previous sections examined price impact and impermanent loss in a two-period setting. In this section, we provide an introductory overview of how these concepts extend to a multi-period framework and discuss their impact on the profits and losses associated with liquidity provision.

We consider a set of periods $t \in \{1,\ldots,T\}$ and a liquidity pool for assets i and j, with asset j being the numeraire. The price $m_{ij} \equiv m$ of the risky asset in the reference market evolves according to the sequence of prices $\{m_t\}_{t=1}^T$, so that the vector of asset prices is $\{P_t\}_{t=1}^T$ where $P_t \equiv (m_t, 1)$. In every period, arbitrageurs trade with the CFMM rebalancing reserves if an arbitrage opportunity arises. If an arbitrage occurs in period $t-1$, the CFMM will enter period t with reserves $R_t = h(P_{t-1})$. Prices and reserves are initialized at $t = 0$ and their values are denoted P_0 and R_0.

As stated in Milionis et al. (2022, 2023), the profits and losses (P&L) of a liquidity provider can be decomposed into the following three components:

$$\mathsf{P\&L}_T = \mathsf{Fee}_T + \mathsf{MRisk}_T - \mathsf{LvR}_T. \tag{39}$$

The first component, Fee_T, accounts for the transaction fees accumulated in the liquidity pool. The second component, MRisk_T, is the market risk exposure. The third component, LvR_T, is a measure of adverse selection costs. Since we discussed fees in Section 7, we now focus on the last two components of (39).

[21] The approximation uses the fact that $(1-\tau)^{-1} = 1 + \tau + O(\tau^2)$.

Loss-versus-Rebalancing

Loss-versus-Rebalancing (LvR) is the period-by-period sum of the impermanent losses defined in (31) and, in the absence of transaction fees, coincides with the arbitrage profits realized over the whole time span:

$$\mathsf{LvR}_T = \sum_{t=1}^{T} L(P_{t-1}, P_t).$$

In LvR, the liquidity provider uses a *rebalancing strategy* as a benchmark: Each time a trade moves the reserves of the risky asset in the CFMM, the liquidity provider replicates the same asset movement in the reference market. That is, she adjusts her reference market holdings of the risky asset $\overline{R}_{i,t}$ so that

$$\overline{R}_{i,t} - \overline{R}_{i,t-1} = \left(h_{i,t} - h_{i,t-1}\right),$$

as illustrated in Figure 10 for $i = 1, j = 2$. This strategy results in the liquidity provider being long in the CFMM and short in the rebalancing portfolio.

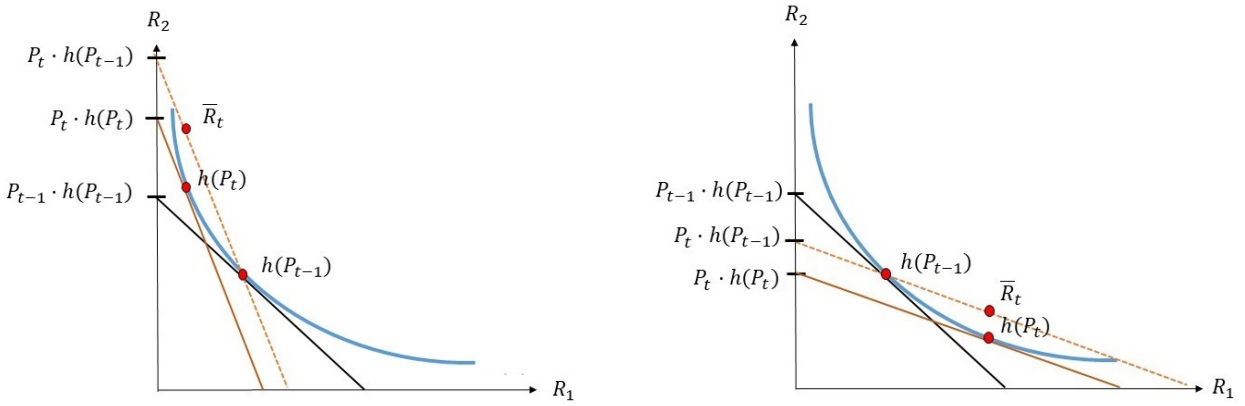

(a) $m_t > m_{t-1}$ (b) $m_t < m_{t-1}$

Figure 10: Rebalancing strategy: the liquidity provider adjusts her holdings in the reference market to offset changes in CFMM reserves.

Rebalancing Strategy as Delta-Hedging for CFMMs

A distinctive feature of the rebalancing strategy is that the liquidity provider is delta-hedging her position, neutralizing the market risk stemming from changes in the price of the risky asset. By replicating in the reference market the opposite asset movements occurring in the CFMM, the liquidity provider effectively creates a market-neutral portfolio. To see this, we

can quantify the market risk component of her portfolio by:

$$\mathsf{MRisk}_T = \sum_{t=1}^{T} \left[h_{i,t-1} - \overline{R}_{i,t-1} \right] \cdot (P_{i,t} - P_{i,t-1}).$$

Under the rebalancing strategy, we have $\overline{R}_{i,t-1} = h_{i,t-1}$, for all t, so the market risk exposure is zero. In this way, by being short in the rebalancing portfolio, the liquidity provider fully hedges the market risk, and her P&L reduces to:

$$\text{Hedged } \mathsf{P\&L}_T = \mathsf{Fee}_T - \mathsf{LvR}_T.$$

This shows that the liquidity provider's profits effectively consist of the fees earned minus the adverse selection costs, as measured by LvR.

Impact of Trading Function Curvature on Profitability

The trade-off linking price impact and impermanent loss to the CFMM curvature, highlighted in Section 6.2, enters $\mathsf{P\&L}_T$ through the determination of Fee_T and LvR_T. The fee income, Fee_T, rises with the quantity of assets exchanged via the CFMM. Hence, Fee_T is positively correlated with CFMM liquidity, which in turn decreases as price impact is reduced. Consequently, a flatter bonding curve, which lowers price impact, leads to higher Fee_T. However, this comes at the expense of increasing the adverse selection component, LvR_T, as lower curvature amplifies impermanent losses. Higher curvature can instead mitigate impermanent losses, reducing LvR_T. In order to maximize the liquidity provider's overall profits, the curvature of the CFMM must be adjusted to achieve an optimal balance between these two opposing effects.

9 Conclusion

We have shown in this paper that standard microeconomic theory sheds a powerful light on the inner workings and optimal design of CFMMs. First, we explained how the main properties of CFMMs, such as their ability to provide reliable oracles for off-chain prices, naturally follow from established results in consumer theory. Transposing the insights of consumer theory enabled us to narrow the design space, for instance by showing that only homothetic trading functions can prevent liquidity providers from scrambling the price signal.

After having established these fundamental results, we used our new tools to gather deeper insights into the economic externalities connecting traders and liquidity providers. Focusing first on traders, we provided explicit conditions under which their costs of using a

CFMM are decreasing in the liquidity of the pool. Then, we turned our attention to liquidity providers, explaining why their impermanent losses are decreasing in the curvature of the bonding curve. Finally, we connected both sides, establishing that the externalities exerted by liquidity providers and by traders are inversely proportional. This finding has important implications for the design of CFMMs since it shows that the preferences of traders are fundamentally divergent from that of liquidity providers: Traders favor CFMMs with low curvature whereas liquidity providers favor CFMMs with high curvature. The challenge for the designer therefore consists in identifying the curvature that strikes an optimal balance between liquidity attraction and fees collection. Solving this mechanism design problem is beyond the scope of the models considered in this paper since it requires specifying the law of motion of prices. We intend to follow this roadmap in future research by leveraging the apparatus laid-out in our paper and extending it to a dynamic setting.

References

Adams, H., Zinsmeister, N., and Robinson, D. (2020). Uniswap v2 core. https://uniswap.org/whitepaper.pdf. Accessed: December 8th 2022.

Andersson, H. (2020). mstable—introducing constant sum bonding curves for tokenised assets. https://medium.com/mstable/introducing-constant-sum-bonding-curves-for-tokenised-assets-6e18879cdc5b. Accessed: August 12 2022.

Angeris, G., Agrawal, A., Evans, A., Chitra, T., and Boyd, S. (2022). Constant function market makers: Multi-asset trades via convex optimization. In *Handbook on Blockchain*, pages 415–444. Springer.

Angeris, G. and Chitra, T. (2020). Improved price oracles: Constant function market makers. In *Proceedings of the 2nd ACM Conference on Advances in Financial Technologies*, pages 80–91.

Angeris, G., Evans, A., and Chitra, T. (2020). When does the tail wag the dog? curvature and market making. *arXiv preprint arXiv:2012.08040*.

Angeris, G., Evans, A., and Chitra, T. (2021a). Replicating market makers. *arXiv preprint arXiv:2103.14769*.

Angeris, G., Kao, H.-T., Chiang, R., Noyes, C., and Chitra, T. (2021b). An analysis of uniswap markets.

Aoyagi, J. (2020). Liquidity provision by automated market makers. *ERN: Other IO: Theory (Topic)*.

Aoyagi, J. and Ito, Y. (2021). Coexisting exchange platforms: Limit order books and automated market makers.

Armstrong, M. and Wright, J. (2007). Two-sided markets, competitive bottlenecks and exclusive contracts. *Economic Theory*, 32:353–380.

Arrow, K. J. and Enthoven, A. C. (1961). Quasi-concave programming. *Econometrica: Journal of the Econometric Society*, pages 779–800.

Bartoletti, M., Chiang, J. H.-y., and Lluch-Lafuente, A. (2021). A theory of automated market makers in defi. In *International Conference on Coordination Languages and Models*, pages 168–187. Springer.

Bergault, P., Bertucci, L., Bouba, D., and Guéant, O. (2022). Automated market makers: Mean-variance analysis of lps payoffs and design of pricing functions. *arXiv preprint arXiv:2212.00336*.

Bichuch, M. and Feinstein, Z. (2022). Axioms for automated market makers: A mathematical framework in fintech and decentralized finance. *arXiv preprint arXiv:2210.01227*.

Capponi, A. and Jia, R. (2021). The adoption of blockchain-based decentralized exchanges. *arXiv preprint arXiv:2103.08842*.

Cartea, Á., Drissi, F., and Monga, M. (2022). Decentralised finance and automated market making: Predictable loss and optimal liquidity provision. *Available at SSRN 4273989*.

Cohen, S., Vidales, M. S., Šiška, D., and Szpruch, Ł. (2023). Inefficiency of cfms: hedging perspective and agent-based simulations. *arXiv preprint arXiv:2302.04345*.

Egorov, M. (2019a). Automatic market-making with dynamic peg. https://classic.curve.fi/files/crypto-pools-paper.pdf. Accessed: March 25 2023.

Egorov, M. (2019b). Stableswap—efficient mechanism for stablecoin liquidity. https://curve.fi/files/stableswap-paper.pdf. Accessed: March 25 2023.

Glosten, L. R. and Milgrom, P. R. (1985). Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of financial economics*, 14(1):71–100.

Goldman, S. M. and Uzawa, H. (1964). A note on separability in demand analysis. *Econometrica: Journal of the Econometric Society*, pages 387–398.

Goyal, M., Ramseyer, G., Goel, A., and Mazières, D. (2022). Finding the right curve: Optimal design of constant function market makers. *arXiv preprint arXiv:2212.03340*.

Jensen, J. R., Pourpouneh, M., Nielsen, K., and Ross, O. (2021). The homogenous properties of automated market makers. *arXiv*.

Lehar, A. and Parlour, C. A. (2021). Decentralized exchanges. *Investments eJournal*.

Martinelli, F. and Mushegian, N. (2019). A non-custodial portfolio manager, liquidity provider, and price sensor. https://balancer.fi/whitepaper.pdf. Accessed: August 12 2022.

Mas-Colell, A., Whinston, M. D., Green, J. R., et al. (1995). *Microeconomic theory*. Oxford university press New York.

Michel, C. (2022). Pricing lp tokens — warp finance hack. https://cmichel.io/pricing-lp-tokens. Accessed: July 30 2022.

Milionis, J., Moallemi, C. C., and Roughgarden, T. (2023). Automated market making and arbitrage profits in the presence of fees.

Milionis, J., Moallemi, C. C., Roughgarden, T., and Zhang, A. L. (2022). Automated market making and loss-versus-rebalancing. *arXiv preprint arXiv:2208.06046*.

Niemerg, A., Robinson, R., and Livnev, L. (2020). Yieldspace: An automated liquidity provider for fixed yield tokens. https://yield.is/YieldSpace.pdf. Accessed: August 12 2022.

Park, A. (2021). The conceptual flaws of constant product automated market making. *Available at SSRN 3805750*.

Port, A. and Tiruviluamala, N. (2022). Mixing constant sum and constant product market makers. *arXiv preprint arXiv:2203.12123*.

Schlegel, J. C., Kwaśnicki, M., and Mamageishvili, A. (2022). Axioms for constant function market makers. *Available at SSRN*.

Tirole, J. and Rochet, J.-C. (2003). Platform competition in two-sided markets. *Journal of the european economic association*, 1(4):990–1029.

Xu, J., Paruch, K., Cousaert, S., and Feng, Y. (2023). Sok: Decentralized exchanges (dex) with automated market maker (amm) protocols. *ACM Computing Surveys*, 55(11):1–50.

A Proofs

Proof of Proposition 3. Let $U(R) = f(u(R))$. Omitting the argument αR in the last of the following equalities, we have that[22]

$$D_\alpha\left(p_{ij}(\alpha R)\right) = D_\alpha\left(\frac{U_{R_i}(\alpha R)}{U_{R_j}(\alpha R)}\right) = \left[\frac{\left(\nabla u_{R_i}(\alpha R)\cdot R\right)u_{R_j}(\alpha R) - \left(\nabla u_{R_j}(\alpha R)\cdot R\right)u_{R_i}(\alpha R)}{u_{R_j}(\alpha R)^2}\right]$$

The last expression does not depend on f_u^2 since it cancels out from the numerator and the denominator. By Euler's Theorem, u 1-homogeneous implies u_{R_i} 0-homogeneous for all i, i.e. $\nabla u_{R_i}(\alpha R)\cdot \alpha R = 0$ for all i, which implies $\nabla u_{R_i}(\alpha R)\cdot R = 0$ for all i given that α is a scalar. It follows that the numerator of $D_\alpha\left(p_{ij}(\alpha R)\right)$ is zero. ■

Proof of Lemma 1. Applying the quotient rule to compute $D_\Delta p\left(\Delta, R\right)$, we have

$$\begin{aligned} D_\Delta p\left(\Delta, R\right) &= D_\Delta\left(\frac{U_{R_i}(R_i-\Delta, R_j+q(\Delta;R))}{U_{R_j}(R_i-\Delta, R_j+q(\Delta;R))}\right) \\ &= \frac{\left(U_{R_iR_j}p - U_{R_iR_i}\right)U_{R_j} - \left(U_{R_jR_j}p - U_{R_jR_i}\right)U_{R_i}}{U_{R_j}^2} \\ &= \frac{\left[U_{R_iR_j}(U_{R_i}/U_{R_j}) - U_{R_iR_i}\right]U_{R_j} - \left[U_{R_jR_j}(U_{R_i}/U_{R_j}) - U_{R_jR_i}\right]U_{R_i}}{U_{R_j}^2}, \end{aligned}$$

with all derivatives evaluated at R^Δ. Multiplying the numerator and the denominator by U_{R_j} yields $s_{ij}(R^\Delta)$ from Eq. (23), which is positive by the definition of quasi-concavity (see Appendix D.1). ■

Proof of Lemma 3. Let us start by proving that the homotheticity of $U(R)$ is sufficient for having 1-homogeneity in $q_{ij}(\Delta, R)$. So assume that U is homothetic. The quantity function is 1-homogeneous if and only if $q_{ij}(\Delta, R) = \Delta p_{ij}(\Delta, R) + R\cdot\nabla_R\, q_{ij}(\Delta, R)$. Rearranging the terms, we obtain

$$R\cdot\nabla_R\, q_{ij}(\Delta, R) = q_{ij}(\Delta, R) - \Delta p_{ij}(\Delta, R). \tag{40}$$

The gradient $\nabla_R\, q_{ij} = (\partial q_{ij}/\partial R_1\ \ldots, \partial q_{ij}/\partial R_N)$ contains the effects on the quantity function

[22] $\nabla u_{R_i}(\cdot)$ denotes the gradient *vector* of the i-th partial derivative of u. This should not be confused with $\nabla u_i(\cdot)$, which is the i-th element of the gradient of u.

of supplying each asset individually. Implicit differentiation of Eq. (4) gives each of these effects:

$$\frac{\partial}{\partial R_k} q_{ij}(\Delta, R) = \frac{U_{R_k}(R) - U_{R_k}\left(R - \Delta\mathbf{i} + q_{ij}(\Delta, R)\mathbf{j}\right)}{U_{R_j}\left(R - \Delta\mathbf{i} + q_{ij}(\Delta, R)\mathbf{j}\right)}. \tag{41}$$

Before combining them, it is important to notice that Eq. (41) is invariant under monotone transformation of U. To see this, let $U = f \circ u$ be a transformation of the trading function u and let $u \equiv u(R), u^\Delta \equiv u\left(R - \Delta\mathbf{i} + q_{ij}(\Delta, R)\mathbf{j}\right)$ and $f^\Delta \equiv f(u^\Delta)$ for brevity:

$$\frac{\partial}{\partial R_k} q_{ij}(\Delta, R) = \frac{\partial}{\partial R_k} \frac{f(u) - f(u^\Delta)}{f(u^\Delta)} = \frac{f_u u_{R_k} - f_u^\Delta u_{R_k}^\Delta}{f_u^\Delta u_{R_j}^\Delta} = \frac{\frac{f_u}{f_u^\Delta} u_{R_k} - u_{R_k}^\Delta}{u_{R_j}^\Delta} = \frac{u_{R_k} - u_{R_k}^\Delta}{u_{R_j}^\Delta} \tag{42}$$

since $f_u / f_u^\Delta = 1$ follows from differentiating both sides of Eq. (4). u produces also the same $q_{ij}(\Delta, R)$ and $p_{ij}(\Delta, R)$ as U since these functions are defined on the bonding curve and are therefore unchanged by a monotone transform.

Now, summation over k of the reserve-weighted effects in Eq. (41), taking u as the 1-homogeneous representation of the trading function, gives

$$R \cdot \nabla_R \, q_{ij}(\Delta, R) = \frac{\nabla u(R) \cdot R - \nabla u\left(R - \Delta\mathbf{i} + q_{ij}(\Delta, R)\mathbf{j}\right) \cdot R}{u_{R_j}\left(R - \Delta\mathbf{i} + q_{ij}(\Delta, R)\mathbf{j}\right)}. \tag{43}$$

By Euler's theorem, the homogeneity of u implies that:

$$\nabla u(R) \cdot R = u(R); \tag{44}$$

$$\nabla u\left(R - \Delta\mathbf{i} + q_{ij}(\Delta, R)\mathbf{j}\right) \cdot \left(R - \Delta\mathbf{i} + q_{ij}(\Delta, R)\mathbf{j}\right) = u\left(R - \Delta\mathbf{i} + q_{ij}(\Delta, R)\mathbf{j}\right) = u(R), \tag{45}$$

as swapping assets keeps the trading function constant. We can then use Eq. (45) to re-express $\nabla u\left(R - \Delta\mathbf{i} + q_{ij}(\Delta, R)\mathbf{j}\right) \cdot R$. Letting $R^\Delta \equiv R - \Delta\mathbf{i} + q_{ij}(\Delta, R)\mathbf{j}$ for brevity, we have that $\nabla u\left(R^\Delta\right) \cdot R^\Delta = \nabla u\left(R^\Delta\right) \cdot R - \Delta \; u_{R_i}(R^\Delta) + q_{ij}(\Delta, R) \; u_{R_j}(R^\Delta)$, using the basis vectors to simplify.[23] Thus,

$$\nabla u\left(R^\Delta\right) \cdot R = u(R) - \left[q_{ij}(\Delta, R) u_{R_j}(R^\Delta) - \Delta \; u_{R_i}(R^\Delta)\right]. \tag{46}$$

Combining Eqs. (44) and (46) to simplify Eq. (43) gives exactly Eq. (40) since $u_{R_j}(R^\Delta)/u_{R_j}(R^\Delta) = 1$ and $u_{R_i}(R^\Delta)/u_{R_j}(R^\Delta) = p_{ij}(\Delta, R)$.

For the necessary part of the lemma, assume that $q(\Delta, R)$ is 1-homogeneous. This is

[23] The basis vectors select the i-th and j-th element of the gradient: $\Delta u_{R_i}(R^\Delta) = \Delta\left[\mathbf{i} \cdot \nabla u(R^\Delta)\right]$, $q_{ij}(\Delta, R) u_{R_j}(R^\Delta) = q_{ij}(\Delta, R)\left[\mathbf{j} \cdot \nabla u(R^\Delta)\right]$.

possible only if the equalities in Eqs. (44) and (45) (and in turn Eq. (46)) are satisfied, which require U to be a monotone transformation of a 1-homogeneous function u. Thus homotheticity is also necessary. ■

Proof of Proposition 5. Corollary 2 and Eq. (23) imply

$$q_{ij}(\Delta, \alpha R) = \int_0^{\Delta} p_{ij}(x, \alpha R)\,\mathrm{d}x = \int_0^{\Delta} p_{ij}\left(\frac{x}{\alpha}, R\right)\mathrm{d}x < \int_0^{\Delta} p_{ij}\left(x, R\right)\mathrm{d}x = q_{ij}(\Delta, R). \blacksquare$$

Proof of Lemma 4. Letting $\bar{P} \equiv \alpha P + (1-\alpha)P'$,

$$\begin{aligned} V(\bar{P}, K) &= \alpha P \cdot h(\bar{P}, K) + (1-\alpha)P' \cdot h(\bar{P}, K) \\ &\geq \alpha P \cdot h(P, K) + (1-\alpha)P' \cdot h(P', K) = \alpha V(P, K) + (1-\alpha)V(P', K). \end{aligned}$$

The inequality holds from the definition of $h(P, K)$ in Eq. (17). ■

Proof of Lemma 5. The separability of $h(\cdot)$ is obvious from the separability of $V(\cdot)$. To prove the latter, we first show the necessary part of the proposition. $U(R)$ being homothetic, there exists a strictly increasing transformation such that $\phi(U(R)) = u(R)$ is homogenous of degree one. Hence

$$\begin{aligned} V(P, K) &= \min_R \{ P \cdot R |\, U(R) \geq K \} = \min_R \{ P \cdot R |\, u(R) \geq \phi(K) \} \\ &= \min_R \left\{ P \cdot R \left| \frac{u(R)}{\phi(K)} \geq 1 \right. \right\} = \phi(K) \min_{R'} \left\{ P \cdot R' \left| u\left(R'\right) \geq 1 \right. \right\} \\ &= \phi(K)\, V(P, 1), \end{aligned}$$

where $R' \equiv R/\phi(K)$ and the penultimate equality uses the homogeneity of $u(R)$.

The sufficient part of the proposition assumes that the portfolio value is separable, i.e. $V(P, K) = \phi(K)\, v(P)$. By definition

$$U(R) = \max \left\{ K |\, P \cdot R \geq \phi(K)\, v(P) \text{ for all } P \in \mathbb{R}^n_+ \right\},$$

hence we have

$$u(R) = \phi(U(R)) = \max \left\{ K' \left| \, P \cdot R \geq K' v(P) \text{ for all } P \in \mathbb{R}^n_+ \right. \right\}$$

with $K' \equiv \phi(K)$. It is clear from the above definition that $u(R)$ is homogenous of degree one and so $U(R)$ is homothetic. ■

Proof of Proposition 8. Differentiation yields $D_{P'}L(P',P) = \nabla_P V(P;K) - \nabla_{P'}V(P';K) = h(P;K) - h(P';K)$. So $L(P',P)$ has critical points at $P' = \alpha P$, $\alpha > 0$, since $D_{P'}L(\alpha P, P) = 0$. Differentiating twice, we obtain $D^2_{P'}L(P',P) = -D^2_{P'}V(P';K)$. Since V is concave in P', $D^2_{P'}E(P';V)$ is negative semi-definite, so $-D^2_{P'}V(P';K)$ is positive semi-definite. $L(P',P)$ is thus convex with global minimum $L(P',P) = 0$ at $P' = \alpha P$ for all $\alpha > 0$. ■

Proof of Proposition 10. We start by showing that the dual representation of the trading set is the portfolio value. $\inf_R \left(P \cdot R - \left[-\delta_{\mathcal{S}(K)}\right](R)\right)$ is clearly achieved at $R \in \mathcal{S}(K)$, giving $[-\delta_{\mathcal{S}(K)}](R) = 0$, as otherwise $[-\delta_{\mathcal{S}(K)}](R) = -\infty$ and so $P \cdot R - [-\delta_{\mathcal{S}(K)}](R)$ would diverge to infinity. Thus, we have

$$\left[-\delta^*_{\mathcal{S}(K)}\right](P) = \inf_{R \in \mathcal{S}(K)} P \cdot R = \min_{R \in \mathcal{S}(K)} P \cdot R = V(P,K). \tag{47}$$

To close the dual representation in the other direction, we take the concave conjugate of $V(P,K) = [-\delta^*_{\mathcal{S}(K)}](P)$, which is

$$\left[-\delta^{**}_{\mathcal{S}(K)}\right](R) = \inf_P \left(P \cdot R - \left[-\delta^*_{\mathcal{S}(K)}\right](P)\right). \tag{48}$$

If $R \in \mathcal{S}(K)$, then the minimum value of reserves that can generate utility K is by definition equal to $\left[-\delta^*_{\mathcal{S}(K)}\right](P) = V(P,K)$, which in turn implies $\left[-\delta^{**}_{\mathcal{S}(K)}\right](R) = 0$. Conversely, if $R \notin \mathcal{S}(K)$, $\left[-\delta^*_{\mathcal{S}(K)}\right](P)$ is equal to $+\infty$, and so $\left[-\delta^{**}_{\mathcal{S}(K)}\right](R) = -\infty$. To summarize:

$$V^*(R,K) = \left[-\delta^{**}_{\mathcal{S}(K)}\right](R) = \begin{cases} 0 \text{ if } R \in S(K) \\ -\infty \text{ otherwise} \end{cases},$$

which is indeed the definition of $[-\delta_{S(K)}](R)$. ■

Proof of Proposition 11. Suppose that the CFMM is initially in equilibrium with reserves R at prices P. Appendix D.2 shows that the inverse function $h^{-1}(R) = (p_{ij}(R), 1)$ is a well-defined bijection under Assumption 1. Then, by the inverse function theorem, we have that

$$s_{ij}(R) = \left.\frac{\partial p_{ij}(R - x\mathbf{i} + q(x,R)\mathbf{j})}{\partial x}\right|_{x=0} = -\left(\left.\frac{\partial h_i(P + y\mathbf{i})}{\partial y}\right|_{y=0}\right)^{-1} = \frac{1}{\ell_{ii}(p_{ij}(R),1)};$$

$$\ell_{ii}(P) = -\left.\frac{\partial h_i(P + y\mathbf{i})}{\partial y}\right|_{y=0} = \left(\left.\frac{\partial p_{ij}(R - x\mathbf{i} + q(x,R)\mathbf{j})}{\partial x}\right|_{x=0}\right)^{-1} = \frac{1}{s_{ij}(h(P))}.$$

The sign change after the second equality of both rows occurs because increasing the relative

price corresponds to *reducing* the Hicksian reserves and vice-versa. ■

B Examples

Example 1. The analogy follows from proving that Hicksian reserves and portfolio value for Uniswap and Balancer are equal.

Uniswap: $U_A(R) = R_1 R_2$. From the optimality condition $P = \lambda \nabla U_A(R)$ we get $R_1 = P_2 R_2$, where asset 1 is our numéraire, so that $P_1 = 1$. Reinserting this solution into the trading function, $R_1 = \sqrt{U_A(R) P_2}$, $R_2 = \sqrt{U_A(R)/P_2}$, and the portfolio value reads $V(P, K_A) = R_1 + P_2 R_2 = 2\sqrt{K_A P_2}$, so that

$$V(P, K_A) = \phi(K_A)\sqrt{P_2}, \text{ where } \phi(K_A) = 2\sqrt{K_A}. \tag{49}$$

Balancer: $U_B(R) = \prod_{i=1}^{n} R_i^{w_i}$. The optimality condition $P = \lambda \nabla U_B(R)$ reads

$$P_i = \lambda \frac{w_i}{R_i} \prod_{j=1}^{n} \left(R_j\right)^{w_j} = \lambda \frac{w_i}{R_i} U(R), \qquad \text{or } \frac{P_i}{P_j} = \frac{w_i/R_i}{w_j/R_j}.$$

For tractability, we focus on the two-asset case. Setting $P_1 = 1$, we get $R_1 = P_2 R_2 w_1/w_2$ and so $U_B(R) = (R_1)^{w_1}(R_2)^{w_2} = \left(\frac{w_1}{w_2} P_2 R_2\right)^{w_1}(R_2)^{w_2}$. It follows that $U_B(R) = \left(\frac{w_1}{w_2} P_2\right)^{w_1}(R_2)^{w_1+w_2}$. Now, solving R_2 as a function of $U_B(R)$ and plugging the result back into the expression of R_1, we get

$$R_2 = \left[U_B(R)\left(\frac{w_2}{w_1 P_2}\right)^{w_1}\right]^{\frac{1}{w_1+w_2}}, \quad R_1 = \left[U_B(R)\left(\frac{w_1}{w_2 P_2}\right)^{w_2}\right]^{\frac{1}{w_1+w_2}}.$$

The portfolio value is therefore given by

$$\begin{aligned} V(P, K_B) = R_1 + P_2 R_2 &= K_B^{\frac{1}{w_1+w_2}} \left[\left(\frac{w_1 P_2}{w_2}\right)^{w_2} + \left(\frac{w_2}{w_1}\right)^{w_1} P_2^{w_2}\right]^{\frac{1}{w_1+w_2}} \\ &= K_B^{\frac{1}{w_1+w_2}} \left[\left(\frac{w_1}{w_2}\right)^{w_2} + \left(\frac{w_2}{w_1}\right)^{w_1}\right] P_2^{\frac{w_2}{w_1+w_2}} \\ &= \phi(K_B) P_2^{\frac{w_2}{w_1+w_2}}, \text{where } \phi(K_B) = \left[\left(\frac{w_1}{w_2}\right)^{w_2} + \left(\frac{w_2}{w_1}\right)^{w_1}\right] K_B^{\frac{1}{w_1+w_2}}. \end{aligned}$$

Now, if we set $w_1 = w_2 = 1/2$, we get

$$V(P, K_B) = \phi(K_B)\sqrt{P_2}, \text{ where } \phi(K_B) = 2K_B.$$

Since $K_B = \sqrt{K_A}$ (see Eq. 49), Uniswap and Balancer with $N = 2$ and $w_1 = w_2 = 1/2$ are indeed equivalent. ■

Example 2. To see that $U(R)$ defined implicitly in Eq. (11) is 1-homogeneous, it is sufficient to prove homogeneity of

$$G(R,U) \equiv \chi(R,U)\left[U^{N-1}\ \sum_{i=1}^{N} R_i - U^N\right]\ + \prod_{i=1}^{N} R_i - \left(\frac{U}{N}\right)^N.$$

$\chi(R,U)$ in Eq. (12) is clearly 0-homogeneous in (R,U). $[U^{N-1}\ \sum_{i=1}^{N} R_i - U^N\ ]$ is N-homogeneous in (R,U), same as $\prod_{i=1}^{N} R_i$ and $(U/N)^N$. So $G(R,U)$ is N-homogeneous in (R,U), meaning $G(\alpha R, \alpha U) = \alpha^N G(R,U)$. Thus, $U(\alpha R) = \alpha U(R)$. ■

Example 3. The concave conjugate for the portfolio value $V(P,K) = 2\sqrt{K(P_1P_2)}$ is

$$V^*(R,K) = \inf_P \left(P_1R_1 + P_2R_2 - 2\sqrt{K(P_1P_2)}\right) = -\sup_P \left(2\sqrt{K(P_1P_2)} - (P_1R_1 + P_2R_2)\right).$$

To compute the conjugate, we need to distinguish two cases. (i) If $R_1R_2 < K$, choose $P_i = \lambda/(2R_i)$ for $\lambda > 0$. Then $2\sqrt{K(P_1P_2)} - (P_1R_1 + P_2R_2) = \lambda\left[\sqrt{\frac{K}{R_1R_2}} - 1\right] > 0$. This expression diverges to infinity with λ, and so $V^*(R,K) = -\infty$. (ii) If $R_1R_2 \geq K$,

$$P_1R_1 + P_2R_2 = \frac{2P_1R_1 + 2P_2R_2}{2} \geq \sqrt{(2P_1R_1)(2P_2R_2)} = 2\sqrt{(R_2R_1)(P_1P_2)} \geq 2\sqrt{K(P_1P_2)},$$

where the first inequality is an application of the AM-GM inequality. We see that $2\sqrt{K(P_1P_2)} - (P_1R_1 + P_2R_2)$ is bounded from above by 0, hence its value is maximized by letting P_1, P_2 and thus $V^*(R,K)$ converge to zero.

To summarize, we have $V^*(R,K) = 0$ if $R_1R_2 \geq K$; $V^*(R,K) = -\infty$ if $R_1R_2 < K$. The conjugate of the portfolio value is indeed equal to the negative characteristic function of the trading set $V^*(R,K) = [-\delta_{\mathcal{S}(K)}](R)$. ■

C Single-Asset Liquidity Provision

We consider a liquidity pool of $N \geq 3$ assets and suppose that the reserves of one of those assets increase. When this occurs, the CFMM should make it cheaper to buy that asset and costlier to sell it. Moreover, trades that do not involve the asset should not be affected by its liquidity provision. An ideal CFMM therefore satisfies these three desiderata of single-asset

liquidity provision:[24]

$$\frac{\partial q_{ij}(\Delta, R)}{\partial R_i} \leq 0, \quad \frac{\partial q_{ij}(\Delta, R)}{\partial R_j} \geq 0, \quad \frac{\partial q_{ij}(\Delta, R)}{\partial R_k} = 0 \text{ for } k \neq i, j, \tag{50}$$

To see which restrictions on the trading function result from Eq. (50), one can implicitly differentiate $q_{ij}(\Delta, R)$ with respect to the reserves of a generic asset k. For an homothetic trading function $U = f \circ u$, where $u(R)$ is 1-homogeneous, we have that

$$\frac{\partial q_{ij}(\Delta, R)}{\partial R_k} = \frac{u_{R_k}(R) - u_{R_k}(R^\Delta)}{u_{R_j}(R^\Delta)} = \frac{1}{u_{R_j}(R^\Delta)} \int_0^\Delta \left[u_{R_k R_i}(R^x) - u_{R_k R_j}(R^x) p_{ij}(x, R) \right] \mathrm{d}x. \tag{51}$$

with $R^x = R - x\mathbf{i} + q_{ij}(x, R)\mathbf{j}$ (also for $x = \Delta$), and $p_{ij}(x, R) = u_{R_i}(R^x)/u_{R_j}(R^x)$. The first equality holds by Eq. (42). The second is an application of the fundamental theorem of calculus. Eq. (51) allows us to achieve the desiderata in Eq. (50) by imposing conditions on the partial derivatives of the 1-homogeneous representation of the trading function. In particular, it is sufficient that u jointly satisfies three properties: (i) diminishing marginal utility: $u_{R_i R_i} \leq 0$; (ii) increasing marginal cross-utility: $u_{R_i R_j} \geq 0$; (iii) separability: $u_{ki}/u_{kj} = u_i/u_j$, which corresponds to $D_{R_k} p_{ij}(\Delta, R) = 0$. Properties (i) and (ii) imply quasi-concavity but not vice versa. Goldman and Uzawa (1964) prove that (iii) holds for all $i \neq j \neq k$ if and only if u is a separable function (and U is its monotone transformation); that is, $u = \sum_{i=1}^N \tilde{u}_i(R_i)$, where each function $\tilde{u}_i$ depends solely on R_i. To summarize, we have shown that:

Proposition 12. *An homothetic CFMM satisfies the desiderata in Eq. (50) if its 1-homogeneous representation u is such that $u_{R_i R_i} \leq 0$ for all i, $u_{R_i R_j} \geq 0$ for all $i \neq j$, and $u_{ki}/u_{kj} = u_i/u_j$ for all $i \neq j \neq k$. Under these conditions, u is quasi-concave and separable.*

D Mathematical Background

D.1 Geometry of quasi-concave functions

We now show that strict quasi-concavity is equivalent to DMRS; i.e. $D_\Delta \mathrm{MRS}_{ij}(R^\Delta) < 0$, where $D_\Delta \mathrm{MRS}_{ij}(R^\Delta)$ is the expression in Eq. (23). Quasi-concavity is what makes negative the sign to the numerator of $D_\Delta \mathrm{MRS}_{ij}(R^\Delta)$ as the denominator, $U_{R_i}(R^\Delta)$, is always positive. Concretely, a twice-continuously differentiable function U is strictly quasi-concave if and only if its Hessian is negative definite along supporting hyperplanes to its level sets (see Fig. 11).

[24] The third desideratum is equivalent to the independence axiom in Schlegel et al. (2022).

That is,[25]

$$x \cdot D^2U \; x < 0 \qquad \text{for all } x \in \left\{ \mathbb{R}^N \setminus \mathbf{0} : \nabla U \cdot x = 0 \right\}. \tag{52}$$

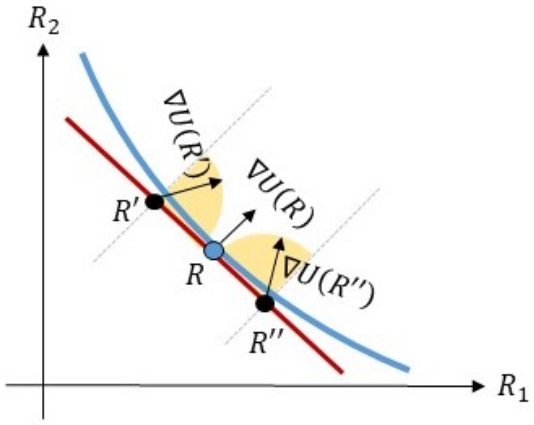

Figure 11: Quasi-concavity

The numerator in Eq. (23) satisfies condition Eq. (52) because it corresponds to a quadratic form of the Hessian D^2U and of the skew gradient $\nabla^\perp U$. The latter rotates the gradient counter-clockwise by 90 degrees, and it is therefore orthogonal to the gradient. Precisely, the numerator of Eq. (52) is given by

$$\nabla^\perp U \cdot D^2U \; \nabla^\perp U, \quad \text{with} \quad \nabla^\perp = \begin{pmatrix} -U_{R_j} \\ U_{R_i} \end{pmatrix} \quad \text{and} \quad D^2U = \begin{pmatrix} U_{R_iR_i} & U_{R_iR_j} \\ U_{R_iR_j} & U_{R_jR_j} \end{pmatrix}. \tag{53}$$

Therefore, by Eq. (52), the numerator in Eq. (23) is negative.

Notice that quasi-concavity defined by Eq. (52) is equivalent to the assumption of positive Gaussian curvature mentioned by Angeris et al. (2020), as the former holds true when $-\left(\nabla^\perp U \cdot D^2U \; \nabla^\perp U\right)/||\nabla U||^3 > 0$ and vice-versa.

D.2 Duality and invertibility of the Hicksian reserves

To demonstrate rigorously and generally the invertibility of Hicksian reserves under Assumption 1, we introduce the concept of superdifferential and supergradient. For a concave function $f \in \mathbb{R}^N$, a *supergradient* is a vector $g \in \mathbb{R}^N$ such that

$$f(y) \leq f(x) + g \cdot (y - x) \text{ for all } x, y. \tag{54}$$

[25] The quadratic form $x \cdot D^2U \; x$ is equivalent to $x^\top D^2Ux$, where $\top$ denotes transposition. The latter expression uses two matrix products; the former uses an inner product followed by a matrix product. We favor the first format to avoid causing confusion with the $\perp$ symbol in Eq. (53).

The supergradient makes the left-hand side of Eq. (54) a global over-estimator of $f(y)$ at each x. The *superdifferential* ∂f is the (closed and convex) set of all supergradients of f. For a singleton superdifferential, we also let ∂f denote its unique supergradient.

For the negative indicator function, we have that

$$\partial[-\delta_{\mathcal{S}(K)}](R) = \left\{P \in \mathbb{R}^N : P \cdot R \leq P \cdot R', \text{ for all } R' \in \mathcal{S}(K)\right\}.$$

Taking R on the bonding curve, it is clear that $\partial[-\delta_{\mathcal{S}(K)}](R)$ gives exactly the prices $P(R, K)$ that minimize the portfolio value at reserves R. So the superdifferential is a singleton set containing the equilibrium prices,

$$\partial[-\delta_{\mathcal{S}(K)}](R) \equiv P(R) = \lambda \cdot \nabla U(R), \tag{55}$$

where λ is pinned-down by the choice of numéraire. Conversely, the supergradient of the dual conjugate $V(P, K) = [-\delta^*_{\mathcal{S}(K)}](P)$ is its gradient, given by the Hicksian reserves:

$$\partial[-\delta^*_{\mathcal{S}(K)}](P) = h(P, K) \tag{56}$$

Combining Eqs. (55) and (56) we can see that conjugate supergradients are inverse functions of each other; that is,

$$\partial\left[-\delta^*_{\mathcal{S}(K)}\right]\left(\partial\left[-\delta_{\mathcal{S}(K)}\right](R)\right) = R, \qquad \partial\left[-\delta_{\mathcal{S}(K)}\right]\left(\partial\left[-\delta^*_{\mathcal{S}(K)}\right](P)\right) = P,$$

thereby establishing the existence of the inverse function

$$h^{-1} : \mathcal{S}^b(K) \to \mathbb{R}^N, \qquad \text{where} \qquad h^{-1}(R) = P(R).$$